# the case for contextualism

Keith DeRose

OXFORD

The Case for Contextualism

This page intentionally left blank

The Case for Contextualism

Knowledge, Skepticism, and Context, Vol. 1

Keith DeRose

Acknowledgements

vi ACKNOWLEDGEMENTS

University of Michigan, Ann Arbor; Massachusetts Institute of Technology; University of Virginia, Charlottesville; Rice University; Syracuse University; Yale University; Rutgers University; University of Vermont; University of Notre Dame; New York University; Brown University; University of Miami; Cornell University; Saint Louis University; and University at Buffalo—SUNY.

I owe special thanks to a few philosophers who have helped me the most. My first job after receiving my Ph.D. was as an assistant professor at the philosophy department at New York University. Happily for me, the environment there was a great and supportive one for doing good philosophical writing, and my philosophical education continued during the three formative years I spent there. This was largely due to my extremely helpful colleagues, John Carroll and Roy Sorensen, who were two of my main teachers during that time. In more recent years, I have profited much from many discussions, and even more e-mail exchanges, with Jon Kvanvig and Jason Stanley.

My greatest philosophical debts are to Rogers Albritton and Peter Unger. It was in discussions with my teacher and dissertation adviser at UCLA, Rogers Albritton, that my thinking on the topics of this book began and began to take shape. I was very fortunate to be among the many philosophers to benefit so much from long discussions with the Socrates of our age. Perhaps my greatest thanks are due to my friend Peter Unger, for his important writings on knowledge and skepticism, from which I learned much; for three years of almost daily philosophical discussions, many of which were on the topics of this book, and almost all of which were enjoyable and enlightening; and for his many comments on various drafts of early papers of mine on these topics. There is a lot of material in here that would or will be very new to Peter and my other old NYU colleagues, John and Roy, and even more in here that would be new to Professor Albritton, were he still alive to read it; my thinking on these matters has come a long way since I worked with them. But it would not have gotten anywhere close to where it is were it not for the jump-start their help gave me early in my philosophical career.

Material from the following papers of mine has been incorporated into this book:

'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* 52 (1992), 913–29.

'Contextualism: An Explanation and Defense', in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology* (Malden, Mass.: Blackwell Publishers, 1999), 187–205.

'Now You Know It, Now You Don't', *Proceedings of the Twentieth World Congress of Philosophy* (Bowling Green, Oh.: Philosophy Documentation Center, 2000), v: *Epistemology*, 91–106.

This page intentionally left blank

Contents

Detailed Contents

This page intentionally left blank

1

Introduction

Contextualism, Invariantism, Skepticism, and What Goes On in Ordinary Conversation

1. Contextualism and the Old Bank Cases

In recent years, the following cases, and other examples like them, have been the focus of a lot of attention from philosophers investigating the meaning of knowledge-attributing and knowledge-denying sentences.[1]

> *Bank Case A.* My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, 'Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays.' I reply, 'No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon.'

---

[1] These cases made their first appearance in the wider philosophical discussion in my (1992: 913), but I had been using them for several years by then. The immediate inspiration for them was Peter Unger's cases involving whether his character John knows that there is milk on the rug in his then just-released (1984: 50–1). Unger contrasted an ordinary ascription of knowledge to John with a case in which philosophical skepticism is being discussed, and the skeptical possibility that 'There is an evil demon deceiving me into false belief in the matter of whether or not there is milk on the rug' is haunting poor John. To make the cases less philosophical and hopefully more convincing, my 'high-standards' Bank Case involved a more moderate hypothesis, the taking seriously of which seemed reasonable given practical, nonphilosophical concerns not present in the 'low-standards' case. These Bank Cases played a prominent role in my dissertation prospectus at UCLA, and were featured later in colloquium talks I gave in the winter of 1989–90, beginning at UCLA in late 1989, and in my dissertation (1990), before appearing in the paper cited at the top of this note. Among others who use such cases in discussions of contextualism, most notable is Stewart Cohen, who employs a very similar case—his Airport Case—in (1999: 58). The material immediately following the case contains a nice, pro-contextualist discussion of it.

Bank Case B. My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, 'Banks do change their hours. Do you know the bank will be open tomorrow?' Remaining as confident as I was before that the bank will be open then, still, I reply, 'Well, no, I don't know. I'd better go in and make sure.'

Assume in each case that the bank will in fact be open on the Saturday in question, and that there is nothing unusual about either case that has not been included in my description of it.

When thinking about Case A, and especially when considering it by itself, without also thinking about Case B, it can seem that

(1) When I say that I know that the bank will be open on Saturday in Case A, my claim is true.

But when considering Case B, especially when not also worrying about Case A, it can seem that

(2) When I admit that I don't know that the bank will be open on Saturday in Case B, my admission is true.

Yet I seem to be in no better position to know in Case A than in Case B. It seems that

(3) If I know that the bank will be open on Saturday in Case A, then I know that it will be open on Saturday in Case B.

So isn't either (1) or (2) wrong? I hope not, because in this book I want to investigate and defend a view according to which both of them—and (3) as well, which we'll discuss in section 4—are true.

Following Peter Unger, we will call the view I will defend 'contextualism': It's a theory according to which the truth-conditions of knowledge-ascribing and knowledge-denying sentences (sentences of the form 'S knows that p' and 'S does not know that p' and related variants of such sentences) vary in certain ways according to the context in which they are uttered. What so varies is the epistemic standard that S must meet (or, in the case of a denial of knowledge,

fail to meet) in order for such a statement to be true. In some contexts, 'S knows that p' requires for its truth that S have a true belief that p and also be in a *very* strong epistemic position with respect to p, while in other contexts, an assertion of the very same sentence may require for its truth, in addition to S's having a true belief that p, only that S meet some lower epistemic standard. In an important sense, the claims in question don't have the same meaning in these different contexts. Making use of David Kaplan's distinction between character and content (1989: 500–7), we can say that the character of 'S knows that p' is roughly, that S has a true belief that p and is in a *good enough* epistemic position with respect to p; this is the element of meaning that remains constant from attribution to attribution. But how good is good enough? That is what varies with context. What context fixes in determining the content of a knowledge attribution is how good an epistemic position the subject must be in to count as knowing that p.

This provides a way for the contextualist to endorse both (1) and (2), even while recognizing that I'm in no better position to know in Case A than in Case B: The standards for what can truly be called 'knowledge' are higher in Case B than they are in Case A. Thus, even while the strength of my epistemic position doesn't vary between the two cases, I meet the lower standards in place in Case A, but don't meet the higher standards in place in Case B. Thus, the truth of both the positive claim to 'know' in A, and the admission that I don't 'know' in B.

At this point, I'm not claiming that we should endorse both (1) and (2). Contextualists often appeal to pairs of cases like these in *arguing* for their view, utilizing the likes of (1) and (2) as premises, and I will pursue such an argument myself in Chapters 2 and 3. But for now, we are just using these cases to *illustrate* what contextualism is by explaining how it handles cases like these, whether or not it is wise to so handle them.

2. Cases Involving Speakers in Different Conversations Talking about the Same Subject

Because he thinks that different standards can constitute the truth-conditions of knowledge ascribing and denying claims on different occasions, the contextualist will allow that one speaker can truthfully say 'S knows that p', while another speaker, in a different context where higher standards are in place, can truthfully say 'S doesn't know that p', though both speakers are talking about the same S and the same p at the same time. To vividly illustrate how this

can happen, at least according to the contextualist, we need cases with very different conversations taking place about the same subject. For that purpose, consider the following example, which begins with a common case set-up, then diverges into two cases, as we follow a couple of characters into very different conversations about the same subject.

Case Set-Up: The Office. Thelma, Louise, and Lena are friends who all work in the same office. Today is their day off, but, before getting an early dinner together, they decide to walk up to the office to pick up their paychecks. Thelma and Lena are also interested in finding out whether a certain colleague is at work, as they are involved in a small office bet with some other workers over whether the often-absent John would show up today. As they pass the door to John's personal office, they see his hat hanging on the hook in hallway, which, in their long experience, has been a sure-fire sign that John is in fact at work. They also hear one working colleague shout to another, 'Why don't you clear that letter with John quick before you send it off?' Satisfied that John is at work and that Thelma and Lena, who bet that he would be, are in a position to collect their winnings from some other office workers, the three friends pick up their checks, go out to dinner together, and then part company, Thelma going to a local tavern to meet other friends, and Louise and Lena each heading in different directions to go home.

Thelma at the Tavern. At the tavern, which is renowned for the low epistemic standards that govern the conversations that take place within its walls, Thelma meets a friend who bet that John wouldn't be at work, and so owes Thelma $2. Thelma says, 'Hey, John was at work today. Pay up, sucker!' When her friend asks, 'How do you know?,' Thelma replies, 'I went up to pick up my paycheck this afternoon. His hat was hanging in the hall outside his office door, and I heard Frank telling someone to quickly check something with John before sending it off.' Satisfied with Thelma's evidence, the friend pays up. Then, wondering whether Lena will know to collect what she is owed by yet another worker, Thelma's friend asks, 'Does Lena know that John was in?' Thelma answers, 'Yes, she was with me. She knows, too.' Meanwhile …

Louise with the Police. Louise has been stopped by the police on her way home. They are conducting an extremely important investigation of some horrible crime, and, in connection with that, are seeking to determine whether John was at work that day. It emerges that they have some reason to think that John was at work and no reason for doubting that (other than the fact that he is often absent from work, which Louise already knows),

but as the matter has become so important to the case, they are seeking to verify that he was there. When the police ask her whether she could testify that John was at work, Louise replies, 'Well, no, I never saw him. I could testify that I saw his hat hanging in the hall, which is a *very* reliable sign that he's at work. And I heard Frank Mercer telling someone to check something with John, as if John were in. But I suppose John could have left his hat on the hook when he went home some previous day. And though it would be a bit strange for Frank not to know whether John was at work, especially that late in the day, I guess he could have been just assuming John was there because John was scheduled to work—and because his hat was in the hall. You should check with Frank. He at least seemed to know that John was in.' When the police ask Louise whether Lena might know that John was in, Louise replies, 'No. She was only at the office very briefly, with me, and didn't see John, either. She has the same reasons I have for thinking John was there, but, like me, she doesn't know that John was there.'

Assume that John was indeed at the office; that the police subsequently verify that with Frank and with a couple of other workers; and that, as luck would have it, Thelma's claim at the tavern that Lena 'knows' that John was at work is made at exactly the same time that Louise tells the police that Lena 'doesn't know' that Frank was at the office. And if you're wondering what Lena is up to at that moment, she is walking home, thinking about the basketball game she plans on watching on the television, not thinking at all about her bet, and not even knowing about any police investigation into any horrible crime.

When thinking about the case set-up and Thelma's conversation at the tavern, and especially when not considering those in connection with Louise's discussion with the police, it can seem that

(4) When Thelma says at the tavern that Lena knows that John was at work, her claim is true.

But when considering the case set-up and Louise's discussion with the police, especially when not worrying about how that discussion relates to Thelma's talk at the tavern, it can seem that

(5) When Louise tells the police that Lena doesn't know that John was at work, her claim is true.

But here we have two speakers talking about the very same subject at the same time. How could it be that Thelma is truthfully saying that Lena 'knows'

that John was at work at the same time that Louise is truthfully saying that Lena 'doesn't know' the very same thing? Here again, contextualism provides a way to reconcile the truth of the two statements: Because different epistemic standards comprise the truth-conditions for Thelma's and Louise's claims, Lena can meet the lower standards governing Thelma's claim, making Thelma's positive claim true, while failing to meet the higher standards for knowledge operative in Louise's context, resulting in the truth of Louise's negative verdict.

3. Contextualism and Invariantism

So, as promised, the story of Thelma and Louise illustrates how, at least according to the contextualist, one speaker can truthfully say 'S knows that p', while another speaker, in a different context where higher standards are in place, can truthfully say 'S doesn't know that p', though both speakers are talking about the same S and the same p at the same time. As with the bank cases, we are not yet arguing for contextualism; we are for now using the cases to illustrate, not defend, the view.

The 'invariantist'—Peter Unger's good name for one who denies contextualism[2]—will not allow that such a thing could happen. According to her, at least given S's situation, there's a single set of standards which, at least as far as truth-conditions go, govern any speaker's assertion about whether or not S 'knows', regardless of these speaker's conversational contexts. Thus, if our two speakers are talking about the same S and the same p at the same time, one saying that S 'knows' that p and the other saying that S 'doesn't know' that p, then even if our speakers are taking part in very different conversations from one another, they can't both be asserting a truth.

In Chapter 2, I will begin to make the case for contextualism. For the rest of this chapter, we will seek to clarify the positions under discussion and set up the rest of the book by briefly explaining the notion of 'strength of epistemic position' that I'm using in my exposition of contextualism; by briefly discussing the issue of just what kind of semantic mechanism might underlie contextualism about 'know(s)'; by distinguishing contextualism, as we're using it, from other nearby positions, some of which also go by the name of 'contextualism'; by distinguishing between the two forms in which contextualism's rival, invariantism, now comes; and by discussing the relation between the ordinary

---

[2] See Unger's (1984), where Unger gives the names 'contextualism' and 'invariantism' the meanings that in many philosophical circles they enjoy today—and the meanings that I am using here.

language basis for contextualism and one of contextualism's most important applications, to the problem of skepticism.

4. 'Strength of Epistemic Position', Comparative Conditionals, and Generic Contextualism

Key to my explanation of contextualism is the notion of *strength of epistemic position*. As I'm here using it, our grasp of that notion is entirely derivative from the concept of knowledge: To be in a strong epistemic position with respect to some proposition one believes is for one's belief in that proposition to have to a high extent the property or properties the having enough of which is what's needed for a true belief to constitute a piece of knowledge.[3] The 'epistemic standards' set by a context are the standards for how strong a position a subject must be in with respect to a proposition, in this very generic sense, for a sentence attributing 'knowledge' to her in the context in question to be true. The contextualist holds that these standards vary according to the speaker's context. The invariantist denies that the standards so vary.

But what properties *are* relevant to whether a subject knows? To answer this would be to move beyond the generic contextualism that I am explaining here, and to start to combine contextualism with substantive views about the nature of knowledge. We will have occasion to do that in Volume II, beginning with that volume's first chapter, where I will combine contextualism with ideas about the nature of knowledge to address epistemological problems, especially the problem of skepticism.

But, before we have any account of the nature of knowledge, there is an intuitive test that I have already hinted at by which we can discern when one subject is at least as well positioned as another.[4] It involves the comparative use of conditionals. Among the grounds one can have for assenting to *If Mugsy is tall, then Wilt is tall* is the comparative knowledge that Wilt is at least as tall as is Mugsy. Given knowledge of the comparative fact, you can know that the conditional is true however the standards for tallness are set, just so long as the

---

[3] Thus, my use of this phrase is very much like one use of 'warrant' popular in recent work in epistemology. Alvin Plantinga (who may be originally responsible for this use of 'warrant,' for all I know), for example, stipulates that by 'warrant', he will mean 'that property or quantity that distinguishes knowledge from mere true belief' (1995: 393). We have to be alert to the possibility that the properties relevant to a true belief's constituting knowledge are a motley crew that have little to do with one another, other than their all being relevant to knowledge.

[4] I have also used this test in (1995: 29–33).

same standards are applied to the two instances of 'tall' in the conditional and so long as those standards are of certain very broad and common types.[5] If you know that Wilt is at least as tall as is Mugsy, then in the relevant circumstances, you're in a position to say: 'Look, I don't know what we're counting as ''tall'' here, but I can tell you this: If Mugsy is tall, then Wilt is tall.' Likewise, where A and B are situations involving S, one can have similarly comparative grounds for assenting to conditionals of the form *If S knows that p in A, then S knows that p in B*. In such a case, the comparative grounds for our assent is our realization that S is in at least as strong an epistemic position with respect to p in situation B as she is in with respect to that same proposition in situation A, and that this conditional seems true on such comparative grounds serves as a good intuitive indication for that comparative fact. Such are our intuitive grounds for the conditional about the Bank Cases I've already presented:

> (3) If I know that the bank will be open on Saturday in Case A, then I know that it will be open on Saturday in Case B.

This conditional seems to be true on comparative grounds such that it should hold no matter how demanding or lax we are being about what we will count as knowledge. It thus serves as a good intuitive indication that I'm in at least as strong an epistemic position in Case A as in Case B. And when we reverse (3), the result is

> (3R) If I know that the bank will be open on Saturday in Case B, then I know that it will be open on Saturday in Case A,

which seems at least as plausible as (3) judged in the relevant comparative way. Together, these conditionals provide a good intuitive indication that I'm in an *equally strong* epistemic position with respect to the bank's being open on Saturday in our two cases: that the differences between the two cases aren't the kinds of differences that are relevant to whether a subject knows something.

The contextualist will hold that some of the differences between the two cases affect the content of my assertions in those cases. However, these differences don't affect the strength of my epistemic position, which is the

---

[5] So long as the standards governing 'tall' are of either what I call (2008) the 'tape measure' or 'group-indexed' varieties, the knowledge that Wilt is at least as tall as is Mugsy suffices for knowing that the conditional is true. Where the standards are of the 'group-relative' variety, the conditional *could* be wrong, but even here, the conditional is verified by Wilt's being at least as tall as is Mugsy, supposing that Wilt and Mugsy belong to the same contextually relevant group. Given that Wilt is at least as tall as is Mugsy, the conditional can go wrong with respect to a single standard for tallness only if the operative standard is a group-relative one, and Wilt and Mugsy belong to different contextually relevant groups.

same in these cases, and thus are not relevant to whether I have the 'knowledge' I claim to have in A, nor to whether I have the 'knowledge' I concede that I lack in B. These matters affect what exactly I'm claiming when I say that I do or do not 'know', but not whether those claims, given those contents, are true.

5. Semantic Mechanism?

One way to go beyond generic epistemic contextualism is to marry one's contextualism to an account of what knowledge is, and to start to specify what's involved in meeting higher or lower epistemic standards. Closely related to this, one can start to explain the relevant mechanisms of conversational dynamics, giving an account of how and when these standards can be changed. Again, I will in those ways be moving beyond generic contextualism—if ever so slightly—in Volume II.

There is another way to move beyond generic contextualism that I won't be engaging in at all. That is to provide an account of the kind of underlying semantic linguistic mechanism responsible for the varying truth-conditions that govern knowledge-attributing and -denying sentences. To explain, it may be best to look at an example of someone who does try to get specific in this way about contextualism.

In his (1996), after quite accurately explaining contextualism in much the same way that I have done here (and elsewhere), and having explained the contextualist's approach to the problem of skepticism, Stephen Schiffer decides to get more specific as he turns to critique contextualism and its approach to skepticism:

The semantic claims about knowledge sentences I've attributed to the Contextualist don't add up to a complete semantic account of those sentences. To what sort of completing semantics might the Contextualist appeal? The important thing is to capture the needed 'indexical' nature of knowledge claims, the fact that different utterances of a knowledge sentence with no apparently indexical terms ('Hillary Rodham Clinton knows that Bill Clinton was a Rhodes Scholar') can express different propositions. I can think of three possible accounts, but none of them will help the Contextualist.   (1996: 325–6)

The attack then comes in three parts. (I respond to the attack in Chapter 5, section 10.) But in discussing Schiffer's attack, I have always just ignored its division into three parts (as have most others I've discussed the matter with), because he basically just makes one argument that gets spread out into multiple

parts, without really utilizing in his argument in any important way anything specific to the different accounts he offers the contextualist, at least so far as I can see.[6]

Immediately following the above quotation, Schiffer explains the first possible contextualist account as follows:

Since a knowledge sentence is supposed to express different propositions in different contexts even if it contains no apparently indexical terms, one naturally thinks of a 'hidden-indexical' theory of knowledge sentences. This would be to assimilate knowledge sentences to sentences like 'It's raining' or 'He's short'. Such sentences express propositions that contain what John Perry has called *unarticulated constituents*: propositional constituents that aren't the semantic values of any terms in the uttered sentence. Thus, an utterance of 'It's raining' might express the proposition that it's raining *in London*, and an utterance of 'He's short' might express the proposition that so-and-so is short *for an NBA center*. This would be the most likely account if we could really take seriously the idea that sentences of the form 'S knows that p' express propositions of the form *S knows that p relative to standards N*. (1996: 326)

But it is not immediately evident what the difference is supposed to be between this 'unarticulated constituent' account and the second option Schiffer offers,[7] which he explains as follows:

A second way of accommodating the alleged context-sensitivity of knowledge sentences would be to claim that the verb 'to know' was itself indexical. Relative to the pretence that the contextualist recognizes only two knowledge standards, Easy and Tough, this would be tantamount to holding that an utterance of 'I know that I have hands' may express, according to the context of utterance, either the proposition that I know$_e$ that I have hands or the proposition that I know$_t$ that I have hands, where the knowing$_e$ relation encapsulates the standard Easy and the knowing$_t$ relation encapsulates the standard Tough. (1996: 327)

---

[6] More precisely, I construe what Schiffer says against the contextualist solution on his first *two* options to be one argument. He has a very different reaction to the contextualist solution on his third way of understanding contextualism, but I don't feel I understand this third option well enough to profitably discuss it. His case against the contextualist solution on this third option depends on, or at least involves, what I take to be a mistaken view about how speakers' reactions to moderate skeptical arguments compare with their responses to the more far-fetched arguments of philosophical skeptics. I take issue with this view at n. 21 of Chapter 5.

[7] After writing this, I found an unpublished paper by Cian Dorr in which he independently professes the same failure to understand this distinction. Concerning belief, rather than knowledge, attributions, Dorr writes: 'Or one might follow Schiffer (1979) and Crimmins and Perry (1989) in positing an "unarticulated constituent" in these sentences as the source of the context-sensitivity: though I confess that I don't really understand the supposed difference between this and the proposal that "believes" is itself context-sensitive.'

Schiffer also proposes a third option for the contextualist at (1996: 327–8), but, as I admit in the previous note, I don't feel I understand this option well enough to profitably discuss it.

These are both accounts that are designed to yield what the contextualist seeks: that the truth-conditions of knowledge sentences will be governed by different epistemic standards in different contexts. (And both accounts yield a wide variety of possible standards. Schiffer only adopts the pretense of two different standards in explaining the second option for ease of exposition. We are to convert his explanation to one where there can be many different standards.) And they're both evidently supposed to move beyond and be more specific than just the generic contextualism we've been working with. And, indeed, *given various theoretical claims* that would give substance to this distinction, there could be real differences between these two accounts in terms of how they predict that 'knows' will 'behave' in various settings. But what such theoretical views does Schiffer suppose here, and what differences in the behavior of 'knows' does he think we should expect given one or the other account? Schiffer doesn't tell us directly. But at least with respect to the first, 'unarticulated constituent' option, we get a hint: We're told that to opt for that model is 'to assimilate knowledge sentences to sentences like "It's raining" or "He's short".' But how, in the relevant ways, do *these* 'behave'? (And what are the relevant ways here?) These are especially pressing questions for one who, like me, thinks that these two examples of Schiffer's behave quite differently *from one another* in ways that might well be relevant at this level of specificity.[8]

---

[8] Here's an example of the kind of difference I have in mind. Suppose speakers in separate examples utter the sentences in question: 'It's raining,' 'He's short.' Suppose these speakers are then asked what they meant, and suppose the questioner provides a menu of options of the type of possible 'unarticulated constituents' that Schiffer would approve of: 'Wha'dya mean "It's raining"? Do you mean that it's raining in London, or that it's raining in Paris, or what?', 'Wha'dya mean "He's short"? Do you mean that he's short for a man, or short for a basketball player, or short for an NBA center, or what?' Now, when the speaker responds to such a question, I think we are *relatively* tolerant, in the case of the speaker who said 'He's short,' of this kind of answer: 'You know, I didn't have anything specific in mind. I was just saying that he was short, and wasn't really thinking about it in any more specific terms than that.' By contrast, it's much *more* puzzling if the speaker who said 'It's raining' replied: 'You know, I didn't have any specific location in mind. I was just saying that it was raining, and wasn't really thinking about it in any more specific terms than that.' I should add that we seem more tolerant about such responses if the speaker of 'It's raining' merely claims not to have thought through the exact extent of the area they were talking about: someone who claims not to have meant anything specific enough to distinguish whether they were talking about London, or the specific area of London they were in at the time, for example. Likewise, it's no big problem if they meant to be referring to their own location—'I meant that it was raining here'—but just hadn't thought through quite how far out they intended their use of 'here' to reach. But it is rather mystifying if they claim not to have any thoughts at all about the location they were talking about. *Comparatively*, it doesn't seem nearly as problematic if a speaker claims not to have had any thoughts at all about the comparison class relative to which their claim that someone is short was made. (Indeed, I think there are cases where gradable adjectives like 'short' are used, but are governed by standards that have nothing to do with comparison classes; see DeRose 2008.) Is this difference good evidence of different semantic mechanisms underlying the context-sensitivity of the two sentences? Who knows? We need a whole lot of theory before we can address that question.

It's also especially pressing if, as I believe, Schiffer has some mistaken views about one of his examples that are very relevant to how it compares with a contextualist account of 'know(s)'. As he hints at the end of the first passage I quote above, Schiffer can't even take seriously the notion that 'sentences of the form "S knows that p" express propositions of the form *S knows that p relative to standards N*'. But he apparently can not only take seriously, but flat-out accepts, the notion that sentences of the form 'S is short' express propositions of the form that S is short for an F, where F is some comparison class supplied by context. Schiffer isn't very forthcoming here about why he thinks the one proposal is hard to even take seriously, while the other is so clearly right that defense (and even explanation) is unnecessary. But, given how some of his related arguments go elsewhere, my *guess* is that this is at least largely because he thinks ordinary, non-philosophical speakers, when asked what they meant by the relevant claims, will be quick to produce an answer of the 'for an F' variety ('I meant that he was short for a basketball player') to a question about the meaning of their use of 'short', but will be less likely to explain themselves in terms of the 'relative to standards N' format when asked what they meant by 'know(s)'.

So it may well be relevant here (though I can't be sure) that Schiffer seems to accept what I call (DeRose 2008) the ' "for an F" myth' about the meaning of gradable adjectives: He seems to think that sentences of the form 'X is A', where A is a gradable adjective, quite generally mean that X is A for an F, where F is some comparison class. At least, he is willing to baldly assert this: 'A sentence of the form "a is tall" expresses a proposition of the form *a is tall for an X*' (1996: 321), without putting a word like 'sometimes' or 'often' in there in front of 'expresses'. But it's also worth noting that this myth is indeed a myth. As I argue in (2008), while the standards that govern applications of gradable adjectives are *sometimes* of the 'for an F' sort—what I call 'group-indexed standards'—that's only one of the types of standards that get used, and on other occasions, gradable adjectives are used relevant to very different standards, so that if you want to say what 'X is A', where A is a gradable adjective, *generally* means, you can't do so as Schiffer does. The best you can do by way of a general account is one according to which it means something like 'X is A relative to standard N'—which, of course, makes it look a *lot* like the account of what 'S knows that p' generally means that Schiffer proposes on behalf of the contextualist, and which makes the general account of the meaning of the relevant sentences diverge in terms of explicit format from the types of answers ordinary speakers are likely to offer when asked what they mean by their uses of gradable adjectives. Would this correction make Schiffer find the 'unarticulated constituent' option hard to even take seriously about

gradable adjectives? Would it make him instead classify gradable adjectives as being indexical terms themselves, whatever exactly he means by that? I just don't know (by any standard!). I'd need to hear a lot more about what Schiffer takes the differences between the categories to be.

At any rate, in this book, I won't be addressing this issue of what kind of semantic mechanism underlies the context-sensitivity of 'know(s)'. I will be engaged here in Volume I with the task of arguing *that* varying epistemic standards govern the truth-conditions of knowledge-attributing sentences, leaving open the matter of *how* (at least in this respect) that all works.

This is not because I think that the second issue is unimportant—or that the two issues cannot interact in important ways. In fact, one way the two issues *might* intersect is very threatening (to say the least) to my view: Perhaps by classifying the various different ways that different context-sensitive terms behave, someone can come up with what they have good grounds to think is a list of the various different ways for terms to be context-sensitive, together with specifications or predictions about how context-sensitive terms of each of the types will 'behave' in various ways. And then it could turn out that 'knows' doesn't fit any of the models, and that therefore we have good grounds for thinking contextualism is false.

At this point, however, I've seen no such argument worth worrying about. One issue such an argument would have to address is whether the list it's based on is incomplete. And, of course, relevant to that is the issue of whether there might be context-sensitive terms—like, perhaps 'know(s)'—that 'behave' in different ways and might be the result of other underlying mechanisms not yet on the list.[9] Here, arguments of the type I will pursue to the effect *that* knowledge-ascribing sentences are context-sensitive can become very relevant to the project of discerning the various *how*'s of context-sensitivity.

6. Which Claims to Take Seriously and the 'Floor' of 'Know(s)'

A question that is often asked of contextualists is: What is the 'floor' for 'know(s)'? How low can the standards get? It's worth briefly discussing this

---

[9] One reason for openness here is that 'know(s)', being of course a verb, can be expected to display behavior much different from the terms it's often compared with—usually adjectives. Exactly how context-sensitive verbs work is a very complex matter about which much openness seems called for at this point. Ludlow (2005) contains some very helpful discussion of this topic.

matter—not because I have particularly well-developed answers to the question, but because addressing this question raises an important methodological issue that I should discuss, however briefly: Which claims to knowledge should we take seriously and seek to accommodate in our theories?

One position that I take seriously (though don't endorse) makes the floor very low indeed: Perhaps in some cases, 'S knows that p' requires for its truth only that S have a true belief that p. It's worth noting that the view that knowledge ascriptions *generally* require only that much for their truth has been seriously defended: see Crispin Sartwell (1991, 1992). I think this cannot possibly be right as a general account of knowledge, but ever since I heard a lecture by Saul Kripke (circa 1985), I've been inclined to accept the contextualist position that I thought was at least suggested there, and that has since been taken by John Hawthorne (2000: 201–6)—namely, that *in certain contexts* the standards for knowledge are so low that *little* (though probably a bit more than nothing) more than true belief is required for the truth of 'S knows that p'. Kripke presented such a case in the above-mentioned lectures: a case in which we're asking 'What did Nixon know, and when did he know it?' and we are unlikely to let Nixon off the hook if his defense is, 'Well, at that point, I merely had true belief. I don't know if I satisfied any of these other requirements for knowledge.' In such a setting, we are likely to seriously describe Nixon as 'knowing' various facts at times even when his claim to being a knower seems to have little more going for it than his having a true belief regarding the proposition in question. Sartwell (1991: 160) and Hawthorne (2000: 202–4) also present some cases, some of them tempting. The contextualist view that true belief is all that's required for the truth of 'S knows that p', at least in some contexts, is easily accommodated (though not necessitated) by the account of knowledge used in my (1995: esp. 36–7), and will further develop in Volume II, where I depict knowledge of p as requiring that p be true and that the subject's belief as to whether p is true match the fact of the matter, not only in the actual world, but in the 'sphere of epistemically relevant worlds' centered on the actual world (the worlds similar enough, in the relevant respects, to the actual world), where the size of the sphere varies with context, being larger in contexts governed by higher standards. The lowest standards, then, may involve a limiting case where the sphere collapses to a point—the actual world—and in such a context, all that's required for knowledge is true belief.

But in cases governed by 'ultra-low' standards of some type, does knowledge really collapse all the way down to true belief? I have my doubts here.[10] Even

---

[10] Hawthorne, too, seems admirably careful not to commit himself to the floor being *that* low, writing, for instance, that 'there are many contexts in which true belief is *just about* all that is required *in*

in such cases, 'S knows that p' may require for its truth *some* further conditions that are easily overlooked (sometimes because they are negative conditions), such as, for instance, that S not have extremely strong reasons for thinking p is false, or that p not be extremely improbable from S's position, in some suitable sense of 'improbable'. It may be because these easily overlooked conditions are satisfied by the subjects involved that speakers in the relevant contexts are inclined to describe these subjects as 'knowers' and that we are inclined to judge such descriptions to be true.

But though I have qualms about whether the cases put forward, even where we agree that the claims about knowledge in them are true, really support the conclusion that knowledge sometimes requires nothing more than mere true belief, I do take at least some of the cases to be ones we should take seriously, where the claims about who 'knows' what are true, and are not to be written off as cases where 'know(s)' is being used non-literally or is being used in some special, separate 'sense' that is best ignored, and I thus agree that the floor is very low, even if perhaps not so low as to allow mere true belief to count.

But then I will be asked: Why stop there? Why, for instance, think that even the truth of p is always necessary for the truth of 'S knows that p'? Indeed, the claim that the truth of p generally fails to be necessary for the truth of 'S knows that p' has also been seriously defended (see Hazlett forthcoming), and there are cases in which speakers quite naturally seem to assert that subjects 'know' various things that aren't true and that are recognized by the speakers not to be true. So, why not accept a contextualized version of this view, denying that the truth of p is always necessary for the truth of 'S knows that p', and hold that in some very lax contexts, subjects can be truthfully described as knowing falsehoods?

In many cases where speakers are inclined to say that subjects 'know' things that the speakers realize are false, intonation provides the signal that these are not to be taken as actual cases of knowledge. In many such cases, the speaker is required to utilize a special intonation when pronouncing 'know(s)'—an intonation that is often indicated in writing by italics or underlining: 'The situation seemed hopeless. Mary *knew* that she wasn't going to survive. Fortunately, she was wrong: The rescue party found her just in time.' In such cases, it would not seem right for the speaker to say the same words 'flat-footedly'—with no special stress or intonation. I take the need for this

---

a positive way' (2000: 204; emphases added) for knowledge. Hawthorne isn't explicit about what else he thinks might be required, but his use of 'in a positive way' would seem to indicate that he has in mind what I'm about to suggest in the text: that there might be some negative conditions (though there's no telling whether the negative conditions I suggest as possibilities would be tempting to him).

special intonation to be a sign that the sentence is not to be taken literally.[11] Note also that such cases go smoothly only when the speaker follows up the strangely stressed statement with an accurate description of the situation that corrects the stressed statement (though sometimes, especially when spinning a very interesting tale, a speaker or writer can appropriately go on for a while before finally letting you in on the real story). Where the speaker realizes that Mary in fact survived her ordeal, it seems very far indeed from correct for the speaker to flat-footedly state that 'Mary knew she wasn't going to survive', and then not to let her listeners in on the fact that Mary actually survived.

These cases seem to be examples of what Richard Holton has called 'protagonist projection' (1997: 626). Here, it seems, the speaker is adopting the perspective of someone—in this case, Mary—and, in the statement in question, describing how the situation appeared from that perspective, rather than providing a true description of the situation.[12]

In some special cases, speakers can do this without special stress or intonation. Especially where speakers are describing what at earlier times (or in other places) was/is very *commonly* thought to be known, they can say, often without unusual stress or intonation, that 'everybody knew' (or that 'we knew' or that 'they knew') something which, as the speaker realizes, was/is in fact false. Here's a nice example I found in the *National Geographic*:

Then in the early 1980s two Australian doctors figured out that peptic ulcers were caused not by stress but most often by bacterial infection. That was a stunning development. Everyone knew that stress caused ulcers! People would say: 'This job is giving me an ulcer.' No, the microbe did it.[13]

The key sentence is presented as an exclamation here, which may function as a sign that something funny is going on, but a similar apparently knowledge-attributing sentence could probably have been presented with no special

---

[11] Note that it's the *need* for intonation that I'm citing: that it seems wrong to assert the sentence without the special intonation. Note also that I'm not claiming that such intonation is quite generally a reliable indicator of non-literal usage. Intonation can be used for all sorts of purposes, so it's unwise to conclude that terms are being used non-literally just because unusually strong intonation is being used. I do think that indicating non-literal usage is the function of intonation in the examples of the use of 'know(s)' under discussion here, but I don't advocate any general account of when it is that intonation serves that particular function. Lacking such a general account, my procedure is to not seek to accommodate any examples where unusually strong intonation seems needed.

[12] Interestingly, one of Holton's own examples involved knowledge:

(8) She knew that he would never let her down, but, like all the others, he did.

Holton writes that 'it would be a rash theorist who argued' on the basis of this example 'that *know* is not factive' (1997: 626).

[13] Joel Achenbach, 'Cat Carrier: Your Cat Could Make You Crazy', *National Geographic*, August 2005.

punctuation: 'Medical discoveries often take us by surprise. Everyone knew that stress caused peptic ulcers—until two Australian doctors figured out in the early 1980s that they were instead most often caused by bacterial infection.' But it is telling that a speaker can get away with such descriptions only in situations in which it's fairly clear that they are projecting themselves into a point of view and describing how the situation looks from that perspective. As others have noted,[14] in similar conversational situations, speakers can also get away with statements like the following: 'Medical discoveries often take us by surprise. For years, peptic ulcers were caused by stress—until two Australian doctors figured out in the early 1980s that they were instead most often caused by bacterial infection.' And surely we don't want to conclude from this that the sentence 'Peptic ulcers were caused by stress' doesn't have as a truth-condition that peptic ulcers were caused by stress!

Compare this with one of the cases of apparently low standards that I *do* take seriously. Let's consider one of Hawthorne's cases:

A clock is stopped and reads '7 P.M.' At 7 P.M., thinking that the clock is working, I set my watch by that clock. My watch now accurately reflects the time, though of course there is no pathway of the sort required by the causal theorists from the fact to the reading. Much later that day (without having a chance to check my watch against other clocks in the interim), I look at my watch, see that it reads '10 P.M.' and think to myself 'It's 10 P.M.; The party begins at 10:30, so it's time to leave for the party.' Someone asks you whether I know that the party begins in half an hour. You are fully apprised of the gettierized etiology of the relevant temporal beliefs. What do you say? Once again, nearly everyone has a strong pull in this context to say 'Yes' and thus to concede that in this case we have gettierized belief that counts as knowledge. (2000: 203)

Here, the speaker is asked whether the subject knows. In this situation, it seems that the speaker can, with perfect propriety, and quite flat-footedly (with no special stress or intonation), inform her questioner that 'John knows that the party begins in half an hour'. In contrast to the cases where a false belief is presented as 'knowledge', this doesn't seem to be a particularly colorful description of the situation about which we might suspect that the speaker is engaging in some special practice like protagonist projection to spice up her description; there is nothing of the 'melodramatic feel' that Holton reports characterize such examples (1997: 627). This seems to be a simple case of a speaker informing her listener of some fact—in this case, that the subject

---

[14] I saw this observation made by Jeremy Fantl in a comment he left on 22 August 2006 to a post put up that same day by Ralph Wedgwood on the weblog, *Certain Doubts*, at <http://fleetwood.baylor.edu/certain_doubts/?p=610>.

knows that the party starts in half an hour. Absent some good reason to think that 'know(s)' is here being used in some special unusual 'sense', or is being used hyperbolically or non-literally—and I don't know of any such good reasons—we should take this use of 'know(s)' seriously. Since the speaker is using the term seriously and properly, and this use isn't based on any false beliefs the speaker has about underlying matters of fact relevant to whether the subject knows, this is a use I think a theory about the meaning of 'know(s)' should try to make come out true, and it's a strike against a theory if it fails to do so. I should add, though, that in this case the subject's belief seems to have considerably more going for it than mere truth: it seems justified, and, depending on how exactly one is imagining the case, there's plenty of room to at least suspect that if the clock had been misrepresenting the time by a large enough margin when the subject saw it, he would not have set his watch by it and/or would have noticed a discrepancy between his watch and the real time before now. So it's not clear just how low a floor this case argues for.

7. Is This Epistemology or Philosophy of Language?

Contextualism and invariantism, as explained above, then, are positions about knowledge attributions (sentences attributing knowledge to a subject) and denials of knowledge—precisely, theses about their truth-conditions. Sadly, and all too commonly, this has been known to give rise to the following type of outburst: 'Your contextualism isn't a theory about *knowledge* at all; it's just a theory about *knowledge attributions*. As such, it's not a piece of epistemology at all, but of the philosophy of language.'[15]

I will address this broad type of charge more thoroughly in Volume II. For now, I will limit myself to just this. First, of course, there's nothing wrong with philosophy of language: It's one of my favorite areas. But what such an objector is usually seeking to deny is not the importance of the philosophy of language, but the importance of this bit of philosophy of language *to epistemology*. And such a denial seems extremely rash. Though contextualism/invariantism is an issue in the philosophy of language, it's a piece of philosophy of language that certainly has the potential to be of profound importance to epistemology. How we should proceed in studying knowledge will be greatly affected by how we

---

[15] These words (especially when the word 'just' is omitted) could be said in a calm spirit of mere classification, in which case this would be no outburst. But I've found such things are often said with a strange degree of hostility—mostly by epistemologists who seem to be sick and fed up with all this stuff they've been seeing and reading about contextualism.

come down on the contextualism/invariantism issue. For contextualism opens up possibilities for dealing with issues and puzzles in epistemology which, of course, must be rejected if invariantism is instead correct. And how could it be otherwise? Those who work on the problem of free will and determinism, for instance, should *of course* be very interested in the issue of what it means to call an action 'free'. If that could mean different things in different contexts, then all sorts of problems could arise from a failure to recognize this shift in meaning. If there is no such shift, then that too will be vital information. In either case, one will want to know what such claims mean. Likewise, it's important in studying knowledge to discern what it means to say someone knows something. If that can mean different things in different contexts, all sorts of problems and mistakes in epistemology, and not just in philosophy of language, will arise from a failure to recognize such shifts in meaning. If, on the other hand, there is no such shift, then we're bound to fall into all sorts of error about knowledge, as well as about 'know(s)', if we think such shifts occur. It's essential to a credible epistemology, as well as to a responsible account of the semantics of the relevant epistemologically important sentences, that what's proposed about knowledge and one's claims about the semantics of 'know(s)' work plausibly together across the rather inconsequential boundary between these two subfields of philosophy.

Here in Volume I, the emphasis will be on philosophy of language, as we seek to adjudicate between contextualism and invariantism. Some epistemology—thoughts about the nature of knowledge, especially—will inevitably sneak in here and there. (For instance, 'intellectualism' about knowledge—a thesis explained below in section 11—will play a very significant role in Chapter 6.) But the focus will be on the philosophy of language—especially as applied to the epistemologically vital terms of 'knowledge'. (I guess if a reasonably accurate label is needed, this volume can be viewed as an exercise in the 'philosophy of epistemological language'.) In Volume II, starting right away in its first chapter, substantive epistemology will begin to play a larger role, as I utilize a very partial but still substantive proto-theory of what knowledge is, together with contextualist semantics, to address a central problem of skepticism. That combined methodology will continue throughout Volume II, as related issues in epistemology are addressed.

8. Contextualism Regarding Other Epistemic Terms

Contextualism, again, is a thesis about knowledge-attributing and -denying sentences. But since there are other terms with tight ties to the concept of

knowledge, we should expect that if contextualism about knowledge is true, there should be corresponding shifts in the content of sentences containing those other terms, especially if the ties between the terms are very tight, and very especially if the other terms are properly analyzed in terms of knowledge. To use David Lewis's words (though he wasn't writing about 'know(s)' when he used them), we should expect the content of knowledge-attributing sentences and the sentences containing the other terms to 'sway together'.[16] For instance, to a first approximation, at least, 'It's possible that p_{ind}'[17] is true if and only if the speaker of the sentence doesn't know that p is false.[18] Given contextualism, then, we should expect that, as the standards for knowledge go up, making it harder for belief to count as knowledge, it should become easier for such statements of 'epistemic possibility' to be true.[19] And, since 'It's certain that p' is the dual of 'It's possible that p_{ind}' ('It's certain that p' is true if and only if 'It's possible that not-p_{ind}' is false), we should expect that as the standards for knowledge go up, making it harder for knowledge attributions to be true, it should also become harder for such expressions of impersonal certainty to be true.[20]

A main theme of Timothy Williamson's justly renowned (2000) is that instead of being viewed as a concept to be analyzed, knowledge should be seen as something useful in the analysis of all sorts of other concepts important to epistemology—and to philosophy of mind as well. If Williamson is right, and if contextualism about knowledge attributions is correct, its potential to spread to other philosophically important terms is great indeed.

---

[16] Lewis (1973: 92); the full sentence, and the sentence which precedes it, which are about the link between similarity and counterfactual conditionals, read: 'I am not one of those philosophers who seek to rest fixed distinctions upon a foundation quite incapable of supporting them. I rather seek to rest an unfixed distinction upon a swaying foundation, claiming that the two sway together rather than independently.'

[17] The subscript 'ind' indicates that the embedded p is to be kept in the indicative mood: very different possibilities are expressed where the p is subjunctive. The subjunctive, 'It's possible that I should not have existed', is just plain good sense, while the indicative 'It's possible that I don't exist' is bizarre.

[18] A more exact analysis of the sentence would still involve the concept of knowledge in a way that, though more complicated, doesn't ruin the point I'm making. See DeRose (1991) for such a more exact analysis.

[19] For more on contextualism and statements of possibility, see my (1990: chapter 2, section V) and my (1998: 69–72).

[20] Though 'It's certain that p' *is* the dual of 'It's possible that p_{ind}', it doesn't mean just that the speaker knows that p, because the provisional analysis of 'It's possible that p_{ind}' as simply being a matter of the speaker not knowing otherwise is just a first approximation, and so isn't exactly correct. For a more exact account of the meaning of 'It's certain that p', see my (1998: 78–9). Expressions of the form 'S is certain that…', where S is a subject—expressions of *personal* certainty—of course, are a very different matter.

We can construe contextualism regarding *justification* as an analogue of what we're calling contextualism about knowledge: According to the contextualist about justification, the standards for justified belief that a subject must meet in order to render true a sentence describing a belief of hers as 'justified' vary with context. The relation between knowledge and justification is controversial, and neither of these forms of contextualism clearly implies the other. If one holds that a belief's being justified is a necessary condition for its being a piece of knowledge, then one may believe that the two contextualisms are closely related: Perhaps it's *because* the standards for justification vary with context that the standards for knowledge so vary. However, it's widely accepted today that more is needed for knowledge than simply justified true belief, and it may be varying requirements for that something more that are reflected in the varying standards for knowledge—in addition to, or instead of, varying standards for justification.

At any rate, in what follows, when I write simply of 'contextualism', I will mean contextualism regarding knowledge attributions; when I mean to refer to contextualism about justification or some other epistemic term, I will explicitly say so.

9. Contextualism is Not a Thesis about the Structure of Knowledge or of Justification

In his influential paper 'A Contextualist Theory of Epistemic Justification' (1978) David Annis presents what he calls 'contextualism' as an alternative to both foundationalism and coherentism in the issue of the structure of justification. Now, I think that even Annis's 'contextualism' is best viewed, not as a structural alternative to those two theories, but rather as a form of foundationalism. But the vital point to be made now is that 'contextualism' as *I* am here construing it is *certainly* not a thesis about the structure of knowledge or of justification. It is, in fact, consistent with both foundationalism and coherentism—and also with various structural options between those two.[21]

If you're a foundationalist, then if you're also a contextualist, you may well come to think of the issue of which beliefs are properly basic (i.e. the issue of which beliefs are warranted to a degree sufficient for knowledge

---

[21] I myself tend to prefer such intermediate options, and defend one, which, following Laurence BonJour, I call 'Weak Foundationalism' (though it should be classified as an alternative to, rather than as a variety of, foundationalism), in my (2005a).

independent of any support they receive from other beliefs) and/or the issue of how strongly supported a belief in the superstructure must be in order to count as knowledge or as a justified belief to be matters that vary according to features of conversational context. And if you're a coherentist, then if you're also a contextualist, you'll probably want to hold that how strongly beliefs must cohere with one another in order to count as knowledge (if they're true), or to count as justified, is a contextually variable matter. So contextualism will certainly color the theories of either structural camp. But contextualism is not itself a structural alternative. Nor does it in any obvious way favor one structural alternative over the other.

10. 'Subject' vs. 'Attributor' Contextualism

Some distinguish between 'subject contextualism' and 'attributor contextualism', or use slightly different titles to mark the same distinction.[22] The basic issue here is whether the truth-value of knowledge attributions is sensitive to features of the context of the putative subject of knowledge, or rather to the context of the attributor—the person describing the subject as a knower or a non-knower. I should be clear that what *I* am calling 'contextualism' is a form of what these folks call 'attributor contextualism'. I don't deny what they call 'subject contextualism'. But it is worth briefly describing what is sometimes called 'subject contextualism' to distinguish it from what I am here calling 'contextualism'.

Some 'subject contextualists' point to examples in which certain features of the surroundings of the putative subject of knowledge which don't constitute any part of his evidence for the belief in question, and which the subject may even be completely oblivious to, affect whether the subject knows. For instance, in Carl Ginet's much-discussed real-barn-among-fakes example,[23] if Henry is driving through a region teeming with fake barns, deceptive enough

---

[22] I've read several papers in manuscript that mark this distinction, some of them with exactly those labels. (When those exact labels are used, their use *may* derive from my distinction between 'subject factors' and 'attributor factors' in part II of DeRose, 'Contextualism and Knowledge Attributions'.) I don't think that 'subject contextualism' has an entirely stable meaning: different philosophers seem to me to use it in slightly different ways, and in recent years, my sense is that it has come to be used mostly to describe only positions of the type we will call 'subject-sensitive invariantism' (see section 11, below). But it seemed to me that there is—or at least was—a use of the term by which it was used to denote the views described in this section, so it is worth distinguishing what I am calling 'contextualism' from such views.

[23] This case plays a prominent role in Alvin Goldman (1976: 772–3); Goldman attributes the example to Carl Ginet (see Goldman 1992: 102 n. 1).

that they would have fooled him if he had come across them, but is luckily encountering the only real barn anywhere about, and he confidently believes that he is seeing a barn, many respond that Henry doesn't know that what he's seeing is a barn.[24] Presumably, though, he would have known it was a barn in a normal situation in which there are no fakes about. Here the abundance of fakes in the region seems to rob the subject of knowledge, even if the subject doesn't know about the fakes. Now, I endorse the intuitions appealed to here: You don't know with the many fakes about, but do know in the normal situation. So, in this way, questions of knowledge are relative to the subject's 'context', where 'context' is being used to describe such extra-evidential features of the subject's situation. So if accepting this is tantamount to being a subject contextualist, then sign me up. But I think it's better—in general, but certainly for our purposes here—to reserve the label 'contextualism' for the more controversial thesis described in section 1, above. For all that's been described here about the position, these 'subject contextualists' are invariantists, at least as we will here be using the relevant terminology.

11. Intellectualism and the Distinction between 'Classical' and 'Subject-Sensitive' Invariantism

Which features of a subject's situation are relevant to whether that subject knows some fact to be the case? The 'subject contextualist' discussed above

---

[24] Some do not share this intuition. In fact, this case doesn't produce in me a strong intuition that Henry doesn't know: See n. 2 of Chapter 2 for some discussion. If you are like me in lacking a strong intuition about this case as it is canonically presented in the philosophical literature, then also join me in changing the case. Imagine that Henry *has* encountered many of the fake barns (let's say 19 of them), and has been fooled by all of them that he has encountered, believing them to be real barns. Now, he's encountering the only real barn in the region (so he is 1-for-20 in his barn judgements), but it seems the same to him as all the other objects he has taken to be barns: He is not in any way more confident of his identification of the object as a barn in the case as compared with all the other instances. ('Yet *another* barn,' he thinks.) Make sure you're imagining the case so that the fakes present an extremely convincing real appearance—just like the real thing—and Henry should have the same visual evidence as he would have had if all the barns around were real, and should be just as justified in thinking that what he is presently looking at is a barn as he would have been if all the barn-like things he had encountered had all been real. Hopefully, though, in this case, in which Henry is not only in a region teeming with fake barns, but has actually been fooled by many of them, and the current, real barn seems to him just like the other objects, you will intuit quite strongly that Henry doesn't know that what he's seeing is a barn. If that doesn't work, imagine that Henry then goes on to encounter 20 more fake barns, confidently taking all of them to be real barns. So now he's 1-for-40 in his barn judgements, with the one 'hit' among the many 'misses' occurring right in the middle of a sea of errors. If you *still* think that Henry knows on the lone occasion on which he's getting it right, then, while there are further modifications to the example we might still try, you are probably beyond the reach of my help here.

deems relevant a factor which some might find a bit surprising. The matter of whether the region he is in is teeming with fake or rather with real barns seems to have no bearing on what Henry's evidence is or whether he is justified in believing that the object he is now viewing is a real barn. Still, it does not seem implausible that this might have an impact on whether Henry knows. This issue does have a bearing on whether Henry's way of telling is reliable in the circumstances he finds himself in: Henry's way of telling is reliable if he's in a normal situation, but not if he's in the Land of Fake Barns. And that seems far from clearly irrelevant to the matter of whether Henry knows. One who takes such matters to be relevant to knowledge can still uphold the highly intuitive thesis that Jason Stanley has helpfully isolated and labeled 'intellectualism', according to which the factors in virtue of which a true belief amounts to knowledge are exclusively truth-relevant, in the sense that they affect how likely it is that the belief is true, either from the point of view of the subject or from a more objective vantage point.[25] For though the actual nature of the other objects in the area that look like barns does not affect how likely it is from Henry's point of view that he's seeing a barn, this matter does affect how likely Henry's belief is to be true from a more objective vantage point.

Recently, however, a number of epistemologists, including Stanley himself, have held that factors about a subject's situation that are not even in this

***

[25] Stanley's official definition of 'intellectualism' makes it the thesis that 'knowledge does not depend upon practical facts' (2005:5). However, in the opening of (2005) where this definition occurs, he is opposing 'practical facts' with factors that he says are '*truth-conducive*, in the sense that their existence makes the belief more likely to be true, either objectively or from the point of view of the subject' (2005: 1). In my construal of intellectualism, the view limits the factors that can affect whether a subject knows to those that are 'truth-relevant' (those that affect how likely it is that the belief is true, either by raising or by lowering that likelihood) instead of Stanley's 'truth-conducive' factors (those that make the belief more likely to be true), because I am construing intellectualism as a thesis about what kinds of factors can be relevant to *whether or not* a subject knows, and factors that operate by making it *less* likely that a subject's belief is true are often knowledge-killers, and so are relevant to that issue. (Stanley phrases things positively, in terms of what factors can be relevant to whether a subject positively knows, and thus focuses only on factors that make the subject's belief *more* likely to be true.) If Stanley is using 'practical' such that any factor that is not truth-relevant is automatically counted as 'practical', then my construal of intellectualism matches his. On the other hand, if he is working with a taxonomy in which the division between 'practical' and truth-relevant factors is *not* supposed to be exhaustive (in which there could be factors that are neither 'practical' nor, in the relevant sense, truth-relevant), then I am here construing 'intellectualism' a bit more narrowly than Stanley does—but harmlessly so. This difference is of no importance in discussing Stanley's own non-intellectualist view, since the non-truth-relevant factors that he takes to be relevant to knowledge are all what he would call 'practical' matters. At any rate, Stanley does take 'intellectualism', as I am construing it here, to be a very widely held assumption in epistemology, and attacking this assumption is a main goal of his book. It is handy to use Stanley's nice label, 'intellectualism', to denote this view.

   Jeremy Fantl and Matthew McGrath label this view 'purism' (2007: 558). Like Stanley, they oppose intellectualism/purism.

broad sense truth-relevant can be relevant to whether that subject knows. The result has been a view that's come to be called 'subject-sensitive invariantism' (SSI, as we will abbreviate it)—though Stanley prefers the term 'interest-relative invariantism'. These deniers of intellectualism hold that whether a subject knows something or not depends in part on such non-truth-relevant, 'practical' matters as the cost (to the subject) of being wrong.

What can be a bit confusing is that many of the non-truth-relevant factors that, according to the various versions of SSI, affect whether a subject knows are the same type of factors that, according to contextualists, affect what knowledge-ascribing proposition gets expressed on a given occasion. In fact, the two views are motivated in part by the desire to get plausible verdicts on whether it would be true for subjects to claim to 'know' things in some of the same types of situations—typically, pairs of situations which differ from one another only in non-truth-relevant factors, but in which it seems correct for a speaker to claim to 'know' in only one of the cases. And, in fact, some seem to think of contextualism as being involved in a denial of intellectualism.

However, as Stanley sees, contextualism typically results from accommodating various intuitions like those elicited by some such examples, while *upholding* intellectualism (2005: esp. 2–3). It is precisely because the contextualist holds fast to the intellectualist assumption that the truth of a given proposition ascribing knowledge to a subject cannot depend on practical, non-truth-relevant matters that he is led to posit that different knowledge relations are denoted in some of the cases in question.

We will discuss SSI, and its comparison to contextualism, in more detail in the chapters to follow—and especially in Chapters 6 and 7. For now, what's important to note is that, though there might be some point in calling SSI a form of 'subject contextualism', in that it makes matters of knowledge turn in part on extra-evidential features of a subject's 'context', SSI is a form of invariantism as we will be using the relevant terms (and is thus well named, at least so far as it is being labeled as a form of invariantism). For, recalling our statement of invariantism from back in section 1 (but now emphasizing a clause key to SSI's counting as a form of invariantism), according to SSI, like other forms of invariantism, *given S's situation*, there's a single set of standards which, at least as far as truth-conditions go, govern any speaker's assertion about whether or not S 'knows', regardless of those speakers' conversational contexts. Thus, if two speakers are talking about the same S and the same p at the same time, one saying that S 'knows' that p and the other saying that S 'doesn't know' that p, then even if our speakers are taking part in very different conversations, they can't both be speaking a truth, according

to SSI. Contextualists, by contrast, insist that both speakers can be asserting a truth.

It's also worth briefly distinguishing this recently hot view, SSI, from the older, intellectualist version of invariantism, which we can call 'classical' invariantism (CI). On CI, regardless of the speaker's conversational situation, *and regardless of the subject's situation*, one set of epistemic standards comprise the relevant truth-conditions that govern all knowledge-attributing and knowledge-denying claims. The truth-relevant factors of the subject's situation, of course, affect whether these invariant standards are met in a given case. Typically, a classical invariantist will hold that what contextualists take to be varying truth-conditions for the relevant claims are in fact just varying conditions for when it would be in some sense *appropriate* or *warranted* to assert the relevant sentences—while the *truth-condition*s of the relevant claims remain fixed through these variations in warranted assertability conditions.

12. A Brief History of Contextualism

Theories according to which there are two senses of 'know(s)'—a 'low', 'weak', or 'ordinary' sense on the one hand, and a 'high', 'strong', or 'philosophical' sense, which is much more demanding, on the other—can be viewed as limiting cases of contextualist views. Such a view was advanced, for instance, by Norman Malcolm in his (1952). For reasons I'll touch on in section 17 below, current contextualist theories don't hold that there are just two different sets of epistemic standards governing the truth-conditions of knowledge attributions, but rather posit such a wide variety of different standards. In looking for precursors to and instances of contextualism here, we will be looking for views that posit such a wide variety of standards.

In important work on knowledge and skepticism in the early and mid 1970s, which culminated in his book *Ignorance: A Case for Scepticism* (1975),[26] Peter Unger argued that, in order to really know something, one must be in a *very* strong epistemic position with respect to that proposition—so strong, in fact, that it would be impossible for anyone ever to be better positioned with respect to any matter than you are now with respect to the matter in question. Though the terminology wasn't in place yet, largely because the contextualist alternative to it wasn't in place yet, what Unger was there defending was *skeptical invariantism*. It was a form of invariantism, and, more

---

[26] This book incorporates, with some improvements, Unger's important journal articles from the early 1970s, while adding new material as well.

particularly, a form of classical invariantism, because, so far as their truth-conditions go, Unger claimed that a single set of epistemic standards govern all attributions of knowledge, in whatever context they were uttered. And it was *skeptical* invariantism because those standards were held to be exceedingly demanding. (Non-skeptical invariantism, then, is invariantism that keeps the standards governing the truth-conditions of knowledge attributions constant, but meetably low.) And Unger drew the skeptical conclusions that were naturally implied by such a stance. Importantly, Unger did admit that varying standards for knowledge govern our use of sentences of the form 'S knows that p', but did not endorse contextualism, because Unger claimed that these varying standards were only standards for whether it was *appropriate* to *say* that S knows; the *truth*-conditions for the sentence, as I've already noted, were, according to Unger, constant, and very demanding. Thus, according to 'early Unger', the skeptic is right when she says we don't know, and we are saying something false (though perhaps appropriate) when, even in ordinary, non-philosophical discussions, we claim to know this or that. This position that we are calling 'classical' invariantism—invariantism about truth-conditions (whether this invariantism is skeptical like Unger's or non-skeptical), combined with variable standards for warranted assertability—is the great rival to contextualism. The 'rival' came first, however: It was partly in response to this 'invariantist' theory of Unger's that the early contextualist views of the mid to late 1970s and early 1980s—like that expressed by David Lewis in a short section of his (1979: 354–5) and in contextualist versions of the relevant alternatives account of knowledge (RA)—were developed.

Most of the prominent examples of contextualism during this time were instances of RA. But RA needn't be a contextualist view, and the contextualist aspect of contextualist versions of RA was not being separated out for explicit separate discussion and critical comparison with invariantism. We will discuss RA, and the conditions under which it is a contextualist view, in the next section.

Later, Barry Stroud, in chapter 2 of his prominent (1984), while not advocating skeptical invariantism, did seek to defend the view by appealing to Unger's idea that the varying standards which can seem to govern the truth-conditions of knowledge attributions might instead just govern their conditions of warranted assertability.

Unger's *Philosophical Relativity* (1984) contained what was at that time—and for some time to come, for that matter—the most complete exposition of the contextualist view. But while this book represented a change of mind for Unger from his skeptical writings of his *Ignorance* period, 'middle Unger' was not advocating contextualism in (1984). Instead, he defended the 'relativist'

conclusion that contextualism and his earlier invariantist views which led to skepticism were equally good theories, and that there simply is no fact of the matter as to which view is correct. To my knowledge, it was here that the terms 'contextualism' and 'invariantism' were first used as we are using them here, to label just these positions, and, more importantly, it was here that these rival positions received their first careful expositions. And in making his case that these two alternatives fought each other to a tie, Unger began the process of identifying the lines of argument that can be used to defend and attack each position. Unger's relativism, defended, as it was, by parity considerations, according to which the advantages and disadvantages of contextualism and invariantism balance each other out in such a way that there is no winner, is a precarious view to defend: Any contextualist who succeeds in defeating invariantism will conquer Unger's relativism as an automatic corollary, and the same will happen for any invariantist who produces a successful argument against contextualism. Thus, it was perhaps no surprise that Unger quickly got off of that fence. With his (1986), 'late Unger' finally cast his lot with contextualism, but did little to counter his earlier arguments that invariantism is the equal of contextualism. Struggling against the invariantist rival that Unger set up remains a main task of contextualism.

Unger then largely left epistemology to pursue work in other areas of philosophy, but thanks mostly to Stewart Cohen, who defended contextualism in a series of important papers, beginning with his (1986) and (1987), the issue was kept alive and kicking.

Interestingly, a nice early statement of contextualism, not usually mentioned when early examples of contextualism are discussed, occurs in a study of ancient philosophy: Gregory Vlastos's (1985) uses contextualism to reconcile the disavowals of knowledge by the Socrates of Plato's early dialogues with the claims to knowledge made by the same Socrates.[27]

---

[27] The actual reconciling work is done by a distinction Vlastos draws between what he calls 'knowledge$_C$' (certain knowledge) and knowledge$_E$ ('elenctic knowledge'), which does not involve the 'fantastically strong standards of knowledge$_C$' (1985: 18), but does require that beliefs be 'elenctically justifiable', as well as true. Thus, through most of the paper, Vlastos seems to be utilizing what looks like some form of the 'two senses of "know(s)"' account, like the account of Malcolm's that I mentioned at the start of this section. (And, interestingly, Vlastos was a colleague of Malcolm's for several years at the Cornell University philosophy department.) However, Vlastos begins the section of his paper in which he draws this distinction with an extremely interesting paragraph that begins with the sentence, 'Let us reflect on our own use of the terms "know" and "knowledge"' (1985: 11), and then proceeds to give a very nice contextualist (though Vlastos doesn't use that word for it) account of the ordinary usage of those terms that seems to express a form of contextualism that's much like the current versions of the view, on which many different standards can govern uses of 'knows', rather than 'two senses of "know(s)"' accounts. Vlastos seems to be thinking of the particular distinction he utilizes throughout much of his paper as a distinction between two of many possible different uses of 'know(s)'. Thus,

My sense is that contextualism grew greatly in notoriety throughout the 1990s, becoming a very well-known and often-studied topic in epistemology. The causes for this, in addition to Cohen's continuing work and my (1992) and (1995), are too many and too complex for me to give anything remotely approaching a good account of them here, but surely prominent among them is the fact that one of the most respected figures in the philosophy of language, and in philosophy generally, David Lewis, wrote and published a prominent contextualist manifesto, 'Elusive Knowledge' (1996).

My general sense is that in the decade of the 2000s, the level of interest in contextualism has remained quite high, but that more of the attention it has been receiving has been in the form of resistance to the view, as new rivals to contextualism, particularly subject-sensitive invariantism and relativism, have arisen, and as general ideas about how context-sensitive pieces of natural language behave have been put forward and applied to epistemic contextualism, often by those who think 'know(s)' doesn't behave much like a context-sensitive term.

13. Contextualism, Invariantism, and Relevant Alternatives

According to the 'relevant alternatives' account of knowledge (RA), which was prominent during the middle and late 1970s and has remained at least somewhat prominent since,[28] the main ingredient which must be added to true belief in order to yield knowledge is that the believer be able to rule out all the relevant alternatives to what she believes. As I indicated above, many of the prominent examples of contextualist theories are instances of RA. But when is a version of RA a contextualist theory?

In one of the standard early presentations of RA, Alvin Goldman (1976) presents various factors which can affect the range of relevant alternatives. These factors may be divided into two groups. First, there are features of the

---

though Vlastos seems to be using a 'two senses' theory through most of his paper, his actual position would seem to be a form of contextualism like those that are currently popular. (Thanks to Gary Ostertag for reminding me of Vlastos's paper in connection with the history of contextualism, and for an interesting e-mail exchange about the paper.)

[28] RA seems to be an especially popular theory of knowledge among non-epistemologists who have to engage in some emergency epistemology in their work. For example, a metaphysician who faces an *epistemological objection* to her views—she's accepting the existence of things such that, according to the objection, we simply cannot know that they exist—is quite likely to utilize RA in dealing with the objection.

putative knower's situation, which I will call 'subject factors'—though please note that by calling them that, I do *not* mean *subjective* (as opposed to objective) factors, but rather factors having to do with the putative *subject* of knowledge and her surroundings (as opposed to the attributor of knowledge). A subject in an ordinary situation can be truly said to know that what he sees up ahead is a barn even if he cannot rule out the possibility that it is just a barn façade. But, Goldman points out, if there are a lot of such façades in the subject's vicinity, then the possibility that what the person is seeing is just a façade *is* a relevant alternative, and the person doesn't know he is seeing a barn, even if what he sees happens to be an actual barn (1976: 772–3).

Second, there are features of the speaker's situation, which I will call 'attributor factors'. Goldman writes, 'It is not only the circumstances of the putative knower's situation, however, that influence the choice of alternatives. The speaker's own linguistic and psychological context are also important.' Goldman suggests that 'if the speaker is in a class where Descartes's evil demon has just been discussed', then certain alternatives may be relevant which ordinarily are not (1976: 776).

Although Goldman draws the distinction between what I am calling subject factors and attributor factors, he does not explain the importance of this distinction. I am stressing it because it is crucial to some of the important claims that some RA theorists have wanted to make about the *meanings* of knowledge attributions and to the closely related matter of distinguishing contextualist from non-contextualist versions of RA.

Those who countenance the possibility of a contextualist version of RA should present RA as a theory about the truth-conditions of knowledge attributions, yielding this rough statement:

> RA: A's claim that 'S knows that p' is true if and only if S has a true belief that p, and S can rule out all the relevant alternatives to p.

This formulation makes room for the possibility of contextualism, but it also allows for the possibility of invariantism, and is neutral between the two. The invariantist can accept the above, in lieu of the simpler, lower-level statement of RA that they'd otherwise gravitate toward—S knows that p iff S has a true belief that p and can rule out all the relevant alternatives to p—by holding that, given a subject's particular situation, the range of alternatives relevant to the truth of *any* speaker's claim about S will be the same, whatever the speaker's context.[29]

___

[29] The invariantist adherent to RA will also want to insist that what counts as 'ruling out' an alternative remains constant no matter what the speaker's context happens to be. Contextualists can

Insofar as an RA theorist allows attributor factors to influence which alternatives are relevant, he is a contextualist. An invariantist can be an RA theorist if he allows only subject factors to influence which alternatives are relevant. Thus, what Goldman calls the 'first' version of RA, according to which 'a complete specification' of the putative knower's situation determines 'a unique set of relevant alternatives' (775–6), is an invariantist version of RA. Goldman does not endorse this view; he says he is 'attracted by the second view' (777), which clearly is a contextualist version of RA.

To illustrate, let us return to the Land of Fake Barns, and consider two situations in which Henry has a good, clear look at what he takes to be—and what in fact is—a barn. In Barn Case A there are no façades around, and all the objects in the vicinity that appear to be barns really are barns, but in Barn Case B the area Henry finds himself in is (unbeknownst to him) teeming with barn façades, some of which he has already encountered, and which have fooled him, although Henry is now luckily looking at the only actual barn in the whole area. RA's explanation for why Henry can be truly said to know in A, but not in B, is that, though Henry cannot rule out the possibility that what he is seeing is a convincing fake barn in either case, the abundance of convincing fake barns in Henry's vicinity makes the possibility that what he is seeing is a fake a relevant alternative in B, while that alternative is irrelevant in Henry's normal surroundings in A. An invariantist can accept this RA account of the difference. Since the difference in whether the fake barn possibility is relevant is due to a difference in subject factors, nothing about this explanation would commit one to saying that whether or not a speaker can truthfully say that Henry 'knows' depends on the speaker's conversational context. One who accepts this RA explanation can hold that, no matter what the speaker's context, any speaker who wants to truthfully say whether or not Henry 'knows' will have to give the positive answer if she's talking about A, but answer negatively if she is issuing a verdict about B.

The *classical* invariantist, recall, is an intellectualist: She believes that whether Henry knows can only depend on broadly truth-relevant factors. Not just invariantists in general, but classical invariantists in particular, can accept the RA explanation for these cases, as it has so far been given. For the difference between these cases that accounts for the difference in whether Henry can be truly said to know, according to this RA explanation, is the abundance of fakes in B as opposed to the absence of fakes in A. And this seems to affect how likely Henry's belief is to be true, not from Henry's own quite limited

hold that both what the range of relevant alternatives is and what counts as 'ruling out' an alternative can vary with context.

perspective (Henry is oblivious to the presence of the fakes in B, and so does not there realize how likely it is that his belief is false), but from some more objective vantage point.

These do not seem to be a pair of cases in which Henry is in equally strong positions to know that what he is seeing is a barn, as is illustrated by the fact that, when judged on comparative grounds (as explained above in section 4), the conditional *If Henry knows in Case A, then he knows in Case B* does not seem to get the comparison right. The classical invariantist holds that how strong an epistemic position a putative knower must be in to be truthfully said to 'know' does not vary from situation to situation. But since, in case pairs like this, the subject is not in equally strong epistemic positions in the two cases, even the classical invariantist can agree that a sentence attributing knowledge to the subject in one case (like our Case A) can be true, while one attributing knowledge to him in another (like our Case B) can be false. And she can use the idea of 'relevant alternatives' to explain the difference. Thus, although most versions of RA allow attributor factors to affect the range of relevant alternatives and are therefore contextualist views, an RA theorist need not be a contextualist, but can be an invariantist, and even a classical invariantist.

By contrast, an invariantist cannot go along with an RA explanation where difference in relevant alternatives is alleged to be due to a difference in attributor factors. Consider two cases where Henry is being discussed, but which don't diverge at all about Henry's situation. Suppose, in particular, that in both cases, Henry is driving through a very normal area, with no fake barns anywhere about. But in these two new cases, there is a big difference between the contexts in which Henry is being discussed. In Barn Case C, the possibility of fake barns has not been raised in the conversation; one speaker is just wondering whether Henry's eyesight is good enough that he could see the barn well enough at that distance to distinguish it from, say, a normal house, and in that context, a second speaker is assuring the first that 'Henry knew it was a barn.' But, following Goldman's lead, suppose that in Barn Case D the speaker is in an epistemology class where fake barns, as well Descartes's evil genius, have just been discussed, and accepted as relevant by all parties to the discussion. Depending on how the classroom discussion goes (and imagine it goes in such a way as to make what follows as plausible as possible), many RA theorists (in particular, those with contextualist inclinations) will think that 'Henry knew it was a barn' is false in Case D due to Henry's inability to rule out the possibility that what he was seeing was just a fake barn together with the relevance of that alternative, while the same sentence is truthfully asserted in Case C, despite Henry's inability to rule out the fake barn alternative,

because of that alternative's irrelevance in C.[30] But invariantists can't go along with this. For here we have just what the invariantist denies can ever happen: We have two speakers each talking about the same subject at the same time, one truthfully saying that the subject 'knows', while the other can truthfully say that the subject does not 'know'.

So, an invariantist can hold to RA if he allows only subject factors to influence what the relevant alternatives are, while an RA theorist who allows attributor factors to affect the range of relevant alternatives will be a contextualist.

Though the contextualist adherent to RA allows that both subject and attributor factors can influence the range of relevant alternatives, in examples like our Barn Cases A and B, she will likely hold that subject factors are the stars of the show: In Barn Case B, factors about Henry's surroundings ensure that the alternative that Henry is seeing a fake will be a relevant alternative in almost any, and quite possibly simply in any, conversational context in which a speaker might find herself talking about Henry in situation B.[31] The speaker might not even know about the fakes, and so might not realize that the possibility of fakes is relevant. Still, the presence of the fakes makes the possibility relevant, no matter the speaker's conversational context. But since everything is normal in Barn Case A, the fakes alternative will be relevant to the truth of claims about that case only in contexts where exceedingly high standards for knowledge are operative. Yet in distinguishing Barn Case C from Barn Case D, attributor factors make the difference. Here, the invariantist will not allow a difference in whether Henry knows, and, consequently, in whether Henry can be truthfully said to 'know', in any context. But the contextualist adherent to RA can hold that the differences in the speakers' contexts in these cases result in the fake barn alternative being relevant to the truth of D's claim,

---

[30] As we will see especially in Chapter 4, contextualists needn't hold that any mention of remote possibilities automatically raises the standards for knowledge to a point where our attributions of knowledge almost all go false. But for now, we will just suppose that what transpired in the epistemology class was sufficient to make 'Henry knew' false, and proceed conditionally on that assumption.

[31] Whether the fake barn alternative is relevant to Henry's knowing in Case B according to *all* allowable standards for knowledge will depend on what the 'floor' is for how low the epistemic standards for knowledge can get. (See section 6, above for some discussion of the issue of possible 'floors'.) If the standards can get so low as to require only true belief for knowledge, then even concerning Case B, the fake barn alternative can be irrelevant. On a contextualist version of RA, standards on which knowledge collapses to mere true belief will be ones on which no alternatives to p are relevant. On a contextualist view that takes the content of a given knowledge attribution to be specifiable in terms of a function from possible subject situations to ranges of relevant alternatives (as described in the last paragraph of section 14, below), standards on which knowledge collapses to mere true belief will involve the function that takes us from all possible subject situations to there being no relevant alternatives.

but not to C's claim, and thus in D's speaker truthfully saying that Henry 'doesn't know' because he can't rule out the fake barn possibility, even while C's speaker truthfully says that Henry 'knows' despite his inability to rule out that possibility.

(But what about first-person claims to know, where the subject and the attributor of knowledge are one and the same? Here, the invariantist will allow only factors that attach to the subject/attributor qua subject of knowledge to affect the range of relevant alternatives, while an RA theorist who allows factors that attach to the person qua attributor of knowledge to affect the range of relevant alternatives will be a contextualist. But dividing the factors that attach to the person in one way from those that attach to her in the other can get extremely messy, and will vary considerably from one theorist to another. Best to distinguish the views by what they say about cases where there isn't such a confusing coincidence of subject and attributor of knowledge.)

14. Against Contextualist Versions of RA That Tie the Content of Knowledge Attributing Claim Directly to What the Range of Relevant Alternatives Is

Thus far, we have been construing RA merely as a theory for explaining why various claims involving knowledge have the truth-*values* they have. But contextualist advocates of RA have often wanted to use the theory to explain what truth-conditions, content, or meaning various claims about knowledge have.

I have been working with what can be thought of as the 'generic' contextualist schema: When the content of knowledge-attributing sentences varies from context to context, what is varying is the epistemic standards, or how strong an epistemic position the subject must be in to count as knowing; the content of a given use of 'S knows that p' is that S (has a true belief that p and) meets the epistemic standards relevant to the context of utterance. This can seem vague and imprecise. One way to try to make the basic contextualist schema more precise utilizes the notion of relevant alternatives. One might think that the content of a given use of a knowledge-attributing sentence can be specified by citing just what range of alternatives are relevant to that use. Thus, the change in content from one use to another will *consist in* a change

in the range of relevant alternatives, and any change in the range of relevant alternatives will amount to a change in the content of the claim. To know by higher standards then would consist in being able to discriminate the truth of the proposition in question from a broader range of relevant alternatives. And just such an idea seemed to be behind some of the moves of some early prominent backers of RA. (Though, historically, what's above in this paragraph gets things backwards: Epistemologists didn't start out with basic contextualism before trying out a way to use relevant alternatives to make it more precise. Rather, contextualism first appeared in the form of contextualist versions of RA, and the contextualist aspects of the view were later abstracted out of it for separate treatment, as contextualism became an item of interest in its own right.) But this attempt to make more precise the basic contextualist schema doesn't work.

Let's start with an example of an early contextualist backer of RA making claims about varying meanings of knowledge-attributing sentences. Gail Stine, in discussing Fred Dretske's example concerning whether a character knows that the animal he is looking at is a zebra (Dretske 1970: 1015−16), writes:

Dretske's zoo example, the animal's being a mule painted to look like a zebra is not a relevant alternative. So what one means when one says that John knows the animal is a zebra, is that he knows it is a zebra, as opposed to a gazelle, an antelope, or other animals one would normally expect to find in a zoo. If, however, being a mule painted to look like a zebra became a relevant alternative, then one would literally mean something different in saying that John knows that the animal is a zebra from what one meant originally and that something else may well be false. (1976: 255)

But here we must be careful. Much depends on *how* the animal's being a painted mule has become a relevant alternative. Unfortunately, Dretske does not discuss variants of his example in which the painted mule possibility *is* relevant, but the most natural, and most secure, way to vary the story to make that possibility relevant is to follow Goldman's lead and make this example analogous to what we're calling Barn Case B. So, suppose that the painted mule possibility has become a relevant alternative because there has been a zebra shortage and many zoos have been using painted mules in an attempt to fool the zoo-going public. If you want, you can imagine that the very zoo that the subject is visiting is so using painted mules, which have actually fooled the subject in the example several times, but that he is now luckily looking at the only genuine zebra in the whole zoo, and that, from his perspective, there's nothing special about this animal, as opposed to all the other animals

that he's just as confidently taken to be zebras.[32] Now, it seems that, though the subject has a true belief that he's seeing a zebra, he doesn't know that he is. On RA, that will be because the possibility that what he's seeing is a cleverly painted mule is, in this case variant, a relevant alternative to its being a zebra. So, here we're imagining that a difference in subject factors—factors about Henry's surroundings—has made the painted mule possibility relevant. This could happen without the speaker knowing it. Would a speaker then *mean* something different by saying that the subject knows that the animal is a zebra? It doesn't seem so, even to me—and I'm one who is of course very open to knowledge attributions having different contents in different situations.

Things seem very different when a difference in which alternatives are relevant is due to a difference in attributor factors concerning the conversational contexts of speakers involved in examples—as between Barn Cases C and D, and between cases analogous to those that could easily be constructed around Dretske's zebra scenario. If an advocate of RA thinks that the attribution of knowledge in C (or a case like it in such a pair) is true and that the denial of knowledge in D (or the case like D) is also true, and thus thinks that there is a difference in the range of relevant alternatives between the cases, then it will seem that there *is* a difference in what the speakers mean by 'know(s)'.

So, respecting those appearances, we will say that attributor factors affect the truth-values of knowledge attributions in a different way than do subject factors: attributor factors working in such a way that they affect the content of the attribution, but subject factors working in a way that does not affect content. These different ways can be explained utilizing the basic, generic contextualist story as follows. Attributor factors set a certain standard the subject must satisfy in order for the knowledge attribution to be true: They affect *how good an epistemic position the subject must be in to count as knowing*. They thereby affect the truth-conditions or the content of the attribution. Subject factors, on the other hand, determine whether or not the subject satisfies the standards that have been set (and, more generally, which standards the subject does and does not satisfy), and thereby can affect the truth-value of the attribution *without* affecting its content: They affect *how good an epistemic position the subject actually is in*.[33]

---

[32] Here I modify Dretske's zebra/painted mules case in a way analogous to how I often modify Ginet's fake barn case; see n. 24, above.

[33] Unger makes a similar division in (1986: 139–40), where he distinguishes between the 'profile of the context', which corresponds roughly to how good a position the subject must be in to count as knowing, and the 'profile of the facts', which corresponds roughly to how good a position the subject actually is in. Unger does not discuss RA at that point, and so does not use this distinction to disentangle

Subject factors, then, are best construed, not as affecting the content or truth-conditions of knowledge attributions, but rather as affecting whether those conditions are satisfied. This severely limits RA's prospects for explaining variations in the content of knowledge attributions. RA, for all I'm saying here, may be a helpful tool for determining or explaining why certain attributions of knowledge have the truth-values they have. So, if your construal of RA is just as a tool for such tasks, I'm not saying anything here against RA as you think of it. Note, however, that for RA to be successful in this capacity, it *must* allow subject factors to affect the range of relevant alternatives, for as Barn Cases A and B clearly show, and as is evident in any case, subject factors can affect these truth-values.

But advocates of RA—particularly several early advocates of the view—have wanted to make claims about the *meaning* of knowledge attributions, thinking that the meaning of knowledge attributions changes from case to case depending on various factors, and thinking that this change in meaning *amounts to* a change in the range of alternatives that are relevant.[34] And we can now see that the content of a given knowledge attribution cannot be specified by citing what range of alternatives is, because that range is a function of subject factors (which don't affect the content of the attribution) as well as attributor factors (which do). There can be a drastic change in the range of relevant alternatives from one attribution to another without there being any change in meaning between the two attributions, then, because the difference in relevant alternatives can be, and often will be, the result of differences in subject factors, which will not have any effect on the meaning of the attribution.[35]

The way to utilize the notion of relevant alternatives in capturing the content of a knowledge attribution in context, it seems, would be to specify that content as a particular function from possible subject situations to ranges

contextualism from RA. He does, however, introduce a complication I'm ignoring here. Unger points out that there are different aspects of knowledge and that in different contexts, we may have different demands regarding various of these aspects. Thus, for example, in one context, we may demand a very high degree of confidence on the subject's part before we will count her as knowing while demanding relatively little by way of her belief being non-accidentally true. In a different context, on the other hand, we may have very stringent standards for non-accidentality but relatively lax standards for subject confidence. As Unger points out then, things may not be as simple as I often make them out to be: Our standards may not be just a matter of how good a position the subject must be in, but rather of how good in which respects.

[34] In addition to the Stine passage we looked at, see, for example Goldman (1976: 775–7, esp. 777), where Goldman seems to think *what proposition is expressed* by a given knowledge attribution can be specified by citing what the range of relevant alternatives is. Something similar seems to be suggested in Lewis (1979: esp. 354–5). Lewis seems to think of the 'conversational score' of a given context, with respect to knowledge attributions and epistemic modal statements, to be something that can be specified by giving the range of possibilities that are relevant in that context.

[35] For some further discussion of this matter, see Brueckner (1994), and my reply, DeRose (1996c).

of relevant alternatives, as I suggest in (DeRose 1996c: 195 n. 6). But, as I also there suggest, having come that far, this question becomes pressing: What's been gained over simply specifying the content of the knowledge attribution as a function from possible subject situations to truth-values of the knowledge attribution, which we easily could have done from the beginning?

15. Against Contrastivism

Jonathan Schaffer (2004, 2005) has defended a contextualist account of know-ledge attributions called 'contrastivism', according to which 'knowledge is a ternary relation of the form Kspq, where q is a contrast proposition' (2004: 77). For a claim of the form 'S knows that p' to be true, on this view, is for S to know that p rather than q, for the contextually relevant contrast proposition q. This view is in many ways similar to contextualist versions of RA—at least similar enough that it runs into the same problems discussed in the above section that RA advocates fall prey to when they tie the content of uses of knowledge-ascribing sentences in context too tightly to the matter of what the relevant alternatives are. The difference is that while possible positions that are recognizably versions of RA (ones that just use RA as a tool for explaining why knowledge attributions have the truth-values they have) don't make the mistake we investigated, Schaffer's contrastivism seems to be explicitly and essentially a theory about the meanings of knowledge attributions that inescapably runs aground on those problems.[36]

   Though the cases discussed in the previous section would suit our purposes fine here, let's use a new example. Suppose Bob and Ernie are both attending a philosophy conference. Unbeknownst to Ernie and the rest of the other conference participants, Bob has brought along to the conference a robot that folks at MIT have constructed that looks and walks just like Bob, and can easily fool even people who know Bob quite well into thinking that it is Bob in situations where only simple behavior is called for. For most of the sessions of the conference, Bob has sent in the robot in his place. The robot has on those occasions walked in, sat down, and seemed to be listening to the paper being presented. As nobody asked the robot any questions, it has fooled everyone at the conference, including Ernie, at each of the nineteen sessions of this long conference that it was at. Everyone has taken the robot to be Bob; nobody even suspects that anything funny is going on. In case others

---

[36] For further critical discussion of Schaffer's contrastivism, see Kvanvig (2007).

at the conference might get suspicious because what they thought was Bob never said anything, Bob himself attended a couple of the sessions instead of the robot, often asking questions and making comments. As luck would have it, Bob himself is attending the current session, but he hasn't yet said anything at this session. Someone asks Ernie if Bob is at the session, and Ernie replies positively. When the person then asks, 'Do you know Bob is here? This is really quite important,' Ernie replies, 'Yes, I know he's here. In fact, that's him sitting right over there in the front row.' And, indeed, it is Bob sitting there: As I've specified, this is one of the few sessions that Bob himself has attended. But Ernie has been fooled by the fakes many times during the conference (indeed, at every session the robot has been at, which has been most of them), and with respect to whether Bob is present, this session doesn't seem any different to Ernie from any of the others.

I take it that if the fake robot has been convincing enough when Ernie has encountered it—and we'll suppose that it has been—Ernie is reasonable in claiming to know that Bob is at the session. Yet it's also quite clear that Ernie's claim to know, though reasonable for him to make, is false. Ernie simply doesn't know that Bob is there, and is speaking falsely (through no fault of his own) when he claims to know it. After all, that robot has been fooling Ernie all conference long in situations that are, from Ernie's point of view, just like this one.

Contrastivism cannot yield both of those clearly correct answers: that Ernie is reasonable in making the claim, and that the claim is nonetheless false. To intuitively yield the result that Ernie's claim is false, the contrastivist must say that the contrast proposition must be or include the possibility that it's an extremely convincing 'fake Bob' robot that's at the session. After all, it's because Ernie can't at this point tell Bob from the robot, together with the robot's real and effective presence at the conference, that Ernie doesn't know that Bob is at the session. But if the possibility of a very convincing 'fake Bob' robot is included in the contrast proposition as part of the very meaning of Ernie's claim to know, then Ernie couldn't reasonably make the claim. He can't reasonably claim to know that it's Bob, as opposed to such a convincing fake, that's at the session, so if what he *meant* by claiming to 'know' that it's Bob includes that he knows that it's Bob rather than such a fake, then his claim would not be at all reasonable.[37]

---

[37] Throughout this argument, I'm assuming that certain values of 'S knows that p rather than q' are what Schaffer would take them to be—here, that Ernie does not know that what he is seeing is Bob rather than a convincing robot. (Generally, Schaffer's judgements about whether a subject knows that one proposition rather than another is the case seem to be guided by something like what we can call this 'contrastive-sensitivity' condition: If S would have believed that p even if q had been the case,

I suppose the contrastivist's best move here—and this is due to a suggestion Schaffer himself made when I asked him about such cases—is to deny that Ernie's claim is false in the example. From the contrastivist's point of view, it can seem that ruling that Ernie's claim is false is to let features of the situation that Ernie is oblivious to problematically affect the meaning of Ernie's claim. It can seem that speakers have more control over their meaning than that, as Schaffer suggested to me. From Schaffer's point of view, all that Ernie is claiming in such a setting in which he is oblivious to the fact that a robot has been impersonating Bob is that he knows that Bob, rather than, say, Fred (another philosopher who has been attending some sessions of the conference), is at the session, and he is not claiming to know that Bob rather than a convincing 'fake Bob' is at the session. But, pre-theoretically, the intuition that Ernie's claim is false is very strong,[38] and, against it, this reply simply assumes

then ~Kspq.) Though Schaffer's judgements about the truth-values of such contrastive-knowledge statements are intuitively plausible, those intuitions seem to me highly suspect, because they conflict with other strong intuitions. Schaffer would judge that

   (a) Ernie doesn't know that Bob, rather than a convincing 'fake-Bob' robot, is at the session,

and (a) does seem intuitively correct. But, in a normal case (in which there is no 'fake-Bob' anywhere on the scene; just a normal conference with normal, human participants), it also seems that

   (b) Ernie knows that Bob is at the session

and, still in a normal case, it finally seems that

   (c) If Ernie doesn't know that Bob, rather than a convincing 'fake-Bob' robot, is at the session, then Ernie doesn't know that Bob is at the session.

Presumably, Schaffer would address this conflict of intuitions about the normal situation by holding fast to (a), and then rejecting, or at least playing around with, (b) and/or (c). (Since (b) and the consequent of (c) are both subject to different interpretations according to Schaffer's contrastivism, there's plenty of room for him to play around with both of them.) However, others might think it best to resolve the conflict by denying (a). (And, for what it's worth, my own inclination is to think that it's (a) that's false about a normal conference situation, *when the sentences are evaluated in accordance with most ordinary standards for knowledge*, and that the reason (a) seems true about such situations is that its inclusion of the robot conjunct makes us tend to evaluate (a) by unusually high standards for knowledge at which it is false. As you can see already from this brief peek at it, the contextualist treatment of this conflict that I'm inclined toward also involves 'playing around with' some of the claims involved, though in a different way than I'm presuming Schaffer would.) And since (a) might well be false about a normal conference situation, it seems quite possible that it is also false about the highly unusual situation we have been considering, in which such a robot actually is present. But since Schaffer's theory won't work anything like he wants it to work unless he's right in his judgements about when contrastive-knowledge statements are true, I am just granting him those judgements here. (Thanks to Jon Kvanvig for some very helpful discussions of these matters.)

   [38] It's important to the intuitive power of this judgement that our case contains the feature that I've stressed in other, structurally similar examples (see nn. 24 and 32): that Ernie actually has been fooled by the robot *many* times, has only been right about whether Bob is present on a couple of occasions, and that on the occasion in question in which Ernie is right, he is not in any way more confident in his judgement than he was on the many occasions in which he was wrong. Also note that Bob has not yet said anything at this session of the conference.

the contrastivist's account of the meaning of the claim, which is precisely the question at issue here. The objector can allow that the speaker is in control of the content of his claim; that only helps the objection. For the speaker is not completely in control of whether the robot possibility is something that Ernie must be able to discriminate from Bob's presence to count as knowing that Bob is at the session: The presence and use of the robot have made that possibility relevant. The contrastivist's account of the content of Ernie's claim is wrong, for that account results in a very counter-intuitive verdict about the truth-value of the claim. What we clearly need is an account of the content of the claim on which that content, *together with features of Ernie's situation*, yield the result that Ernie needs to be able to discriminate Bob from the robot to be able to truthfully claim to know that Bob is at the session. We have good reason to reject a view like contrastivism, according to which the content of Ernie's claim to knowledge *all by itself* determines whether Ernie has to be able to discriminate Bob from the robot to be able to truthfully claim to know, no matter how the subject factors of his situation are arrayed.[39]

16. The Contextualist Approach to Skepticism and to What Goes on in Ordinary Conversation

Contextualist theories of knowledge attributions have almost invariably been developed with an eye toward providing some kind of answer to philosophical skepticism. For many of the most powerful skeptical arguments threaten to show that we just plain don't know what we ordinarily think we know. They thus threaten to establish the startling result that we never, or almost never, truthfully ascribe knowledge to ourselves or to other mere mortals.

But, according to typical contextualist analysis, the skeptic, in presenting her argument, manipulates the semantic standards for knowledge, thereby at least threatening to create a context in which she can *truthfully* say that we know nothing or very little. To the extent that she succeeds in so raising the standards, we correctly sense that we could only falsely claim to know such things as that we have hands. Why then are we puzzled? Why don't we simply

---

[39] Analogous to my suggestion for RA in the last paragraph of section 11, above, the remedy would seem to be to loosen up the connection between the content of the knowledge claim and what's included in the contrast proposition, and identify the content of the knowledge claim with a function from possible subject situations to contrast propositions. But that serious complication of the view would seem to take away from it a lot of the nice features Schaffer claims for it.

accept the skeptic's conclusion and henceforth refrain from ascribing such knowledge to ourselves or others? Because, the contextualist continues, we also realize this: As soon as we find ourselves in more ordinary conversational contexts, it will not only be true for us to claim to know the very things that the skeptic now denies we know, but it will also be wrong for us to deny that we know these things. But then, isn't the skeptic's present denial equally false? And wouldn't it be equally true for us now, in the skeptic's presence, to claim to know?

What we fail to realize, according to the contextualist solution, is that the denials that we know various things that are governed by the standards that the skeptic at least threatens to install are perfectly compatible with our ordinary claims to know those very propositions. Once we realize this, we can see how both skeptical denials of knowledge and our ordinary attributions of knowledge can be correct.

Thus, it is hoped, our ordinary claims to know can be safeguarded from the apparently powerful attack of the skeptic, while, at the same time, the persuasiveness of the skeptical argument is explained. For the fact that the skeptic can, and perhaps sometimes even does, install very high standards which we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in more ordinary conversations and debates.

The success of such an approach to skepticism hinges largely on the contextualist's ability to explain how the skeptic raises the standards for knowledge in the presentation of her argument. If such an explanation can be successfully constructed, such a solution to skepticism can be very attractive, at least to some, and can provide a powerful motivation for accepting contextualism. So applying contextual analysis to the problem of skepticism is the main topic of the second volume I am planning.

But while philosophical skepticism has drawn much of the attention of con-textualists, support for contextualism should also—and perhaps primarily—be looked for in how 'know(s)' is utilized in non-philosophical conversation. For as the cases we've already looked at here illustrate, we do seem to apply 'know(s)' differently in different contexts, a phenomenon that, at least on the surface, seems to promise significant support for contextualism. And the contextualist's appeal to varying standards for knowledge in his solution to skepticism would rightly seem unmotivated and ad hoc if we didn't have independent reason to think such shifts in the content of knowledge attributions occur.

Additionally, the contextualist solution to the skeptical puzzles might fail to do justice to the impact the skeptical arguments can have on us if it were only in the presence of such arguments that the standards for knowledge were shifted upward. Why? Well, it is a fairly natural reaction to the skeptical

arguments to suppose that they induce us to raise the standards for knowledge. In most classes of any size in which I've first presented skeptical arguments to introductory students, some of them will pursue such an analysis of the argument's force. Usually, they propose a version of the 'Two Senses of "Knows"' theory I mentioned at the start of section 12. Students sometimes label the two senses they posit 'weak' and 'strong', or sometimes 'low' and 'high', and once 'regular' and 'high octane'. But though many will suspect that the skeptic is somehow 'changing the subject' on us, she certainly isn't doing so in any very obvious way—as is shown by the fact that some students in most introductory classes will reject the suggestion that any such thing is going on. Yet, if there were just two senses of 'know(s)'—one normal and quite common, and the other very strong and brought on only in contexts in which philosophical skepticism is being discussed—it would probably be quite clear to us that the skeptic was doing something fairly new and different when she started using 'know(s)' in the 'high octane' sense, and it would probably be pretty obvious she was 'changing the subject' on us, and the arguments wouldn't seem to be a threat to our ordinary knowledge. On the other hand, if, as current contextualists hold, the standards for knowledge vary with context even in ordinary, non-philosophical conversations, and the skeptic is utilizing mechanisms for the raising of epistemic standards that we're familiar with from such ordinary conversations, then it would seem much more likely that the skeptic's argument would strike us as threatening our knowledge ordinarily so called, since the skeptic's use of 'know(s)' would much more likely pass for what ordinarily goes on with the use of the term. As Gail Stine wisely wrote:

It is an essential characteristic of our concept of knowledge that tighter criteria are appropriate in different contexts. It is one thing in a street encounter, another in a classroom, another in a court of law—and who is to say it cannot be another in a philosophical discussion? … We can point out that some philosophers are very perverse in their standards (by some extreme standard, there is some reason to think there is an evil genius, after all)—but we cannot legitimately go so far as to say that their perversity has stretched the concept of knowledge out of all recognition—in fact they have played on an essential feature of the concept.   (1976: 254)

17. Relativism, Fervent Invariantism, and the Plan for this Volume

In the present volume, then, I will seek to make the case for contextualism by appeal to what goes on in ordinary, non-philosophical discussion, before

turning, in Volume II, to the application of contextualism to the problem of philosophical skepticism. In this volume, I will be arguing for the superiority of contextualism over invariantism, in both of the latter's main varieties—classical invariantism and subject-sensitive invariantism.

There are a couple of rivals to contextualism that I will *not* be arguing against here, though my defense of contextualism may be relevant to the evaluation of these rivals to it. In each case, there are good critical discussions of these rivals that others have put forward, and I don't have anything of significant value to add to these discussions, beyond what is already contained in my defense of contextualism here.

First, in recent important work, John MacFarlane (2005) has advanced a 'relativist' account of the semantics of knowledge-ascribing sentences as an alternative to both contextualism and invariantism. Here is how MacFarlane describes his view:

I propose a semantics for 'know' that combines the explanatory virtues of contextualism and invariantism. Like the contextualist, I take the extension of 'know' to be sensitive to contextually determined standards. But where the contextualist takes the relevant standards to be those in play as the context of *use*, I take them to be those in play at the context of *assessment*: the context in which one is assessing a particular use of a sentence for truth or falsity.   (2005: 197)

This is part of a larger movement in which MacFarlane and other philosophers and linguists have advocated relativist semantics for many different terms in natural language. Viewed from certain angles, the proposed relativist semantics are quite radical, and one of the tasks MacFarlane sets for himself is defending the coherence of his relativist approach (2005: 225–30). Assessing the coherence and plausibility of relativist semantics, both generally and in the specific case of knowledge-ascribing sentences, is a large and important topic that is attracting a lot of attention these days.[40] What useful contribution I'm in a position to make to this discussion right now is in answering certain objections to contextualism. MacFarlane presents his relativism as arising from the ashes of the non-relativist views of contextualism and invariantism. It is because he thinks both of these non-relativist views run aground of important considerations that MacFarlane is motivated to try his radical hypothesis. In the case of contextualism, the relevant objections are of the type I take on in Chapters 4 and 5 of this volume.

---

[40] Jason Stanley has a nice, and quite critical, discussion of relativism concerning knowledge attributions in chapter 7 of Stanley (2005). The literature on relativist semantics in general is large, and growing too quickly for me to keep up with. But one critical piece that I found very valuable, and found myself largely in agreement with, is Richard Heck's (2006), which focuses on MacFarlane's relativist account of statements about 'future contingents', but raises important general points about relativism in the process.

In defending contextualism from these objections (and especially by showing, in Chapter 4, how contextualism can handle phenomena of disagreement, an important motivating consideration for many relativists), I hope to remove the motivation for relativism—at least in the case of knowledge attributions, though my defense of epistemic contextualism will also have application to contextualist accounts of other terms.

Second, in recent work, culminating in their (2005), Herman Cappelen and Ernie Lepore have advanced a view—which we might well call 'fervent invariantism'—on which the list of genuinely context-sensitive terms is shockingly short. 'Know(s)' doesn't make their list. So, of course, Cappelen and Lepore are among my opponents. And they are particularly troublesome opponents for me because a key to many of my responses to objections to epistemic contextualism is to appeal to the 'behavior' of what I take to be clearly context-sensitive terms (as you will see, 'tall' is my favorite example), and Cappelen and Lepore deny that even these terms are context-sensitive. So, when I began work on this book (this was before Cappelen and Lepore 2005 was published, but I was working on its precursors), I projected a good deal of material battling against Cappelen and Lepore: arguing that the tests they propose for context-sensitivity, depending on how they are understood, are either hopelessly bad tests, or are actually passed by the terms that Cappelen and Lepore mean to exclude, or else fail to allow in even the terms on their shockingly short list, and in no case do they provide for a good argument for the authors' fervent invariantism. But the need for my attack on Cappelen and Lepore has since been obviated by Hawthorne (2006), which rather compactly (and successfully, in my view) takes on Cappelen and Lepore along grounds similar enough to my own projected attack to render the latter obsolete. So on this score, I just refer the reader to Hawthorne (2006)—but then also to the fervent invariantists' response to Hawthorne (Cappelen and Lepore 2006: 473−81).

Briefly, then, here is what will follow in the rest of this volume. Chapters 2 and 3 concern the positive case for contextualism. Chapter 2 makes and explains the main argument for contextualism, the argument from the ordinary usage of 'know(s)'. Chapter 3 explores the connections between knowledge and assertion in order to answer the main source of resistance to the argument from ordinary usage, to otherwise bolster that argument from Chapter 2, and also to generate another positive argument for contextualism. Chapters 4−7 answer objections to contextualism, and also further develop and clarify the view in ways relevant to the objections being considered. Chapter 4 concerns the phenomena of disagreement, which some suspect are troublesome for contextualism, and how disagreement can be approached from a contextualist

perspective.[41] Chapter 5 answers a variety of objections to contextualism built on how 'know(s)' behaves in comparative judgements of content and within metalinguistic claims and speech reports and belief reports. Chapters 6 and 7 both wrestle with considerations which can appear to threaten both contextualism and subject-sensitive invariantism. By answering these concerns on behalf of contextualism and arguing that SSI does face serious objections in these areas, I make the comparative case for contextualism over SSI. Chapter 6 concerns how the two views fare with respect to considerations tied up with an 'intellectualist' approach to knowledge. Chapter 7 is about how well the two views handle the connections that knowledge at least seems to bear to appropriate assertion and rational action. It has been suggested, especially by Hawthorne, that SSI has important advantages over contextualism here. But I will counter the objections to contextualism that have emerged on this front, and then go on to argue that contextualism fares better than does SSI on this front.

---

[41] An interesting and important literature has developed lately on 'the epistemology of disagreement'—roughly, on epistemological issues that are raised by cases where intelligent, well-informed people disagree with us: how should such disagreement affect our beliefs, and how does such disagreement affect the epistemic status of our beliefs? See, e.g., Kelly 2005, Feldman 2006, Christensen 2007, Elga 2007, and the essays in Feldman and Warfield forthcoming. I will not be attempting to make any contribution to that important literature here. (Though I have seen and heard of attempts to use contextualism to address the issues mentioned above, I am not convinced that contextualism can play a significant role in solving those issues, and am convinced that it cannot play anything approaching a dominant role.) What I will instead be addressing is the objection to contextualism laid out in section 2 of Chapter 4: That in certain situations, contextualism rules that speakers are not contradicting one another, when it intuitively seems fairly clear that they are doing just that.

2

The Ordinary Language Basis for Contextualism

1. The Main Argument for Contextualism

The best grounds for accepting contextualism come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, non-philosophical talk: What ordinary speakers will count as 'knowledge' in some non-philosophical contexts they will deny is such in others. As J. L. Austin observed, in many ordinary settings we are easy, and say things such as (Austin's example), 'I know he is in, because his hat is in the hall.' But, even with no philosophers in sight, at other times speakers get tough and will not claim to know that the owner was present based on the same evidence; as Austin notes, 'The presence of the hat, which would serve as proof of the owner's presence in many circumstances, could only through laxity be adduced as a proof in a court of law' (1979a: 108). We needn't invoke anything as unusual as a high-stakes court case to find such variation—as I'm sure Austin realized. A wide variety of different standards for knowledge are actually used in different ordinary contexts.

Following Austin's lead, the contextualist will appeal to pairs of cases that forcefully display this variability: A 'low-standards' case (henceforth 'LOW') in which a speaker seems quite appropriately and truthfully to ascribe knowledge to a subject will be paired with a 'high-standards' case (henceforth, 'HIGH') in which another speaker in a quite different and more demanding context seems with equal propriety and truth to say that the same subject (or a similarly positioned subject) does not know.[1] When it is handy to refer to particular

---

[1] The contextualist follows Austin's lead in using such pairs of cases, but note first that the contextualist, as I have presented him, goes beyond Austin here in having the speaker go so far as to deny knowledge in HIGH, while Austin merely notes that it would be wrong in HIGH to cite as a proof the meager grounds that sufficed in LOW. As we will see below, this is a crucial difference. And note secondly that though the contextualist follows Austin's lead in appealing to such pairs of

cases, the case pairs we looked at in Chapter 1 will function as our main examples. (And in one respect I followed Austin's lead very closely indeed: some of the details of the office case are directly inspired by Austin's example.) In Chapter 1, we were using those examples to explain contextualism; here we will investigate the argument for contextualism that they can generate.

The case from such considerations is the contextualist's main argument, and will be our focus in this chapter. This type of basis in ordinary language not only provides the best grounds we have for accepting contextualism concerning knowledge attributions, but, I believe, is evidence of the very best type one can have for concluding that *any* piece of ordinary language has context-sensitive truth-conditions. In this chapter, we'll look in more detail at some of the features of the ordinary use of 'know(s)' that make for a compelling case for the contextualist account of that verb (and its cognates), and explain why these features militate so strongly for a contextualist treatment. So we'll be looking both at what the data from ordinary language are, and at the methodology that takes us from the data to a contextualist conclusion. The case will be completed in Chapter 3, where I will battle the main source of resistance to the argument: the charge that what the contextualist takes to be varying truth-conditions for knowledge-attributing sentences is really just a variation in the conditions under which those sentences can be appropriately asserted.

cases, I don't think we can cite Austin as himself being a contextualist. As I mentioned in Chapter 1, opponents of contextualism characteristically accept the evident facts about typical and, in the relevant sense, appropriate linguistic behavior that such case pairs display. What is controversial is whether these varying standards for when ordinary speakers *will* attribute knowledge and when they're *warranted* in attributing knowledge reflect varying standards for when it is or would be *true* for them to attribute knowledge. (The main burden of Chapter 3 will be to wrestle against the invariantist suggestion that what the contextualist takes to be a variation in the truth-conditions of knowledge-attributing sentences is really only a variation in their conditions of warranted assertability.) The real controversy, then, is over a crucial issue about which Austin seems (unwisely, I believe) quite unconcerned. In (1979b), for example, attached to the (to my thinking) wonderful passage, 'Certainly, then, ordinary language is *not* the last word: in principle it can everywhere be supplemented and improved upon and superseded. Only remember, it is the *first* word,' Austin drops this (to my thinking) troubling footnote: 'And forget, for once and for a while, that other curious question "Is it true?" May we?' (1979b: 185). As H. P. Grice points out, in connection with this troubling footnote, Austin 'consistently avoids the words "true" and "false," using instead such expressions as "it will not do to say" …' and seeks 'to put on one side the question whether such applications would be inadmissible because they were false, because they would lack a truth-value, or for some other reason' (1989: 13). Austin displays a similar attitude in his epistemological writings, studiously avoiding issues of whether our epistemological claims (especially claims about what is and is not 'known') are true or false. Thus, though Austin is something of a granddaddy of contextualism, in that he put on display some of the linguistic phenomena contextualists appeal to, he himself should not be classified as a contextualist. All grant that in some way our use of 'knows' is governed by different standards in different contexts. The question is: Are the varying standards we can discern those that govern the truth-conditions of the sentences, or merely their warranted assertability conditions? Insofar as I can discern a clear answer to this in Austin, it is (a peculiarly aggressive), 'I don't care.'

2. Mutually Reinforcing Strands of Evidence

The contextualist argument based on such pairs of cases ultimately rests on the key premises that the positive attribution of knowledge in LOW is true, and that the denial of knowledge in HIGH is true. Why think that both of these claims are true?

Well, first, and most directly, where the contextualist's cases are well chosen, those are the fairly strong intuitions about the cases, at least where each case is considered individually. Here we appeal to how we, competent speakers, intuitively evaluate the truth-values of particular claims that are made (or are imagined to have been made) in particular situations.[2] Our intuitions about

[2] It's important to note that the intuitions being appealed to concern the truth-values of the claims made in the examples. They are intuitions that would be elicited by the question 'Is the speaker's claim true?', and not by the question, 'Does the subject know?' If I were presented with the contextualist's cases from Chapter 1, and then asked, 'Do the subjects know [the propositions in question (that the bank will be open on Saturday, that John was at work)]?', what intuitions I would have about that matter would be far more wavering and uncertain than are my intuitions that the claims made in the cases are true, and I believe what intuitions I would have about the object-level question would be largely derived from my intuitions that the claims in the examples are true: Because I find the claims attributing 'knowledge' to the subject in the low-standards cases to be true, I would be led to think the subjects in those cases do know; and because I would judge the claims denying 'knowledge' to the subject in the high-standards cases to be true, I would be led to think the subjects in those cases do not know. (Here, I'm assuming that I'm presented with the cases separately, and so don't consider them together. If I were presented with the high- and low-standards cases together, then the pressure to give the same verdict about whether the subject in question knows in the two cases would be great—and greater than is the pressure to rule that one or the other of the claims made within the cases (that the subject 'knows' in LOW, and doesn't 'know' in HIGH) must be false.)

In recent years, there has been a lot of skepticism expressed about the use in epistemology of intuitions concerning whether or not subjects in particular cases know propositions to be true. For a very nice overview of this important debate over methodology, see Nagel (2007). I am often a bit skeptical of many such uses myself. To give just one example, it is very widely accepted in epistemology today that the subject (Henry) doesn't know that he's seeing a barn in Ginet's real-barn-among-fakes example (that we considered in Chapter 1). And I must admit that, perhaps moved by the authoritative judgements of others, after I encountered the example, I rather uncritically joined this consensus for a few years, and was happy to rest arguments on the claim that Henry didn't know. However, when Ruth Millikan first confronted me with skepticism about that judgement, I had to admit that it really was not all that clearly correct. (See Millikan 1984: 329–30. I encountered Millikan's skepticism in person a few years later, in the winter of 1989–90, at which time she referred me to her paper.) Key to Millikan's doubts here is the fact that, at least in the canonical presentation of the example, Goldman (1976: 772–3), though there are many convincing fake barns in Henry's vicinity, Henry has not encountered any of them. Though I still was—and still am—personally inclined to judge that Henry doesn't know in the case as canonically described, since encountering Millikan's skepticism, I have become skeptical enough myself that I haven't taken this judgement to be clearly enough correct for it to be usable as the premise of a good argument. (Fortunately, the purpose for which I was using the example in that discussion with Millikan was still served when I modified it in the way described in n. 24 of Chapter 1. The judgement that Henry does not know that he's seeing a barn in this *modified* version of the example does seem solid enough to make it usable as a premise of fairly convincing arguments. Is it *certainly* right? Of course not. But it seems pretty solid to me, so far as premises of

such matters can be wrong, of course, but still, they are among our best guides when evaluating semantic theories, especially when we are careful to avoid the types of situations where we are likely to be misled—and there is good reason to think that the intuitions the contextualist appeals to here are among those we should invest much confidence in, as we will see.

To reinforce these intuitions, the contextualist can appeal to the facts that his cases display how speakers in fact, and with propriety, use the claims in question. As we can sense, perhaps with even more certainty than that with which we judge the truth-values of the claims, speakers do in fact use 'know(s)' in the way described, and appropriately so: They will in fact, and with apparent propriety, ascribe 'knowledge' in situations like LOW, yet will deny 'knowledge' when they find themselves in conversational circumstances like HIGH. This supports the premises that both of the imagined claims are true, since generally (though there are some exceptions), one cannot properly claim something that from one's own point of view (given one's beliefs about the underlying matters of fact³ relevant to the claims in question) is false. So, since the contextualist's cases do not involve speakers who are involved in

philosophical arguments go.) And in general, I have since become wary of the intuitions used in much epistemology.

However, at least with regard to many cases, the intuition about whether a claim involving 'know(s)' made in the examples is true can be at least fairly strong, and when that intuition can be buttressed in the ways we are about to investigate, I think it can become the basis of a strong philosophical argument.

³ The basic idea behind this use of 'underlying matters of fact' is that these are the types of features of situations on which speakers base their claims that subjects do or do not 'know' the propositions in question, and which a philosopher would specify in describing a case to a respondent whose intuitions the philosopher is seeking about whether the subject knows something or about whether a speaker's claim that the subject 'knows' is true. Whether p is true, for instance, is something that is very relevant to whether a subject in a case knows that p or whether a speaker's claim that the subject 'knows' is true. Usually, when describing a case to a respondent when we seek the respondent's intuition as to whether the subject knows that p or whether a speaker's claim that the subject 'knows' is true, we will somehow specify whether p is true. If we fail to specify this, and the respondent asks us whether p is true in the example, this typically is information we should provide, and so the matter of whether p is true typically is a relevant 'underlying matter of fact'. But the matters of what 'know(s)' means, whether a contextualist or an invariantist semantics for 'know(s)' is correct, and under what conditions someone knows something to be the case, are *not* 'underlying matters of fact'. We shouldn't specify these matters in describing a case to a respondent, and if the respondent asks us to specify these matters, we should refuse to do so, explaining that these are matters under dispute, which we're hoping might be adjudicated by the use of the intuitions we are soliciting. Our hope is that the respondents' implicit grasp of what knowledge is and what 'know(s)' means might guide them, based on their grasp of what the underlying facts of the situation are, to a more-or-less secure judgement about whether the subject knows and/or whether a claim that the subject 'knows' is true. (In some cases where it's controversial whether the embedded p in 'S knows that p' is true, we might ask respondents whether a given subject knows that p in an attempt to determine what the relevant intuitions are concerning the truth-value of p as well as whether 'S knows that p' is true. In such a case, the truth of p will not be an 'underlying matter of fact'. We won't specify whether p is true, and

some mistaken belief about a relevant underlying matter of fact, there is good reason to think that their claims, which are made with perfect propriety, are true, and it's a bad strike against a semantic theory if it rules these claims to be false, as it seems invariantism will have to rule with respect to one or the other of the claims.

To cite these facts concerning common and appropriate uses of 'know(s)' as evidence for the truth of our imagined speakers' claims that is thoroughly independent of our intuitions that their claims are true would be to engage in some illegitimate double-accounting. For it is no doubt largely because we already sense that the speakers in the cases are using 'know(s)' in very common and appropriate ways that we intuitively judge that their claims are true.

But it is possible for claims to appear to be true without being in the relevant sense appropriate, though our intuitions that such claims are true will often be more tentative in such a case. Indeed, if I can be permitted to engage briefly in some personal history, my earliest attempts at formulating a pair of cases to support contextualism suffered from just such a defect. My adviser, Rogers Albritton, objected, as near as I can remember, 'Nobody would really talk that way!' I replied that it didn't matter whether people would talk that way. All I needed was that such a claim would be true, and that certainly was my judgement about the truth-value of the claim. He would have none of that, and answered, quite sternly, 'Look, if you're going to do ordinary language philosophy—and that's what you're doing here—you'd better do it right.' So I returned to the cases, and adjusted them so that what I had the speakers say seemed more natural and appropriate, as well as true. And, to my surprise, the changes did not only make the argument read more nicely, which was the only pay-off I was expecting, but actually made it more compelling. Albritton never explained to me why the examples should be constructed so that what's said is natural and appropriate beyond insisting that that's how ordinary language philosophy should be done. (He seemed to think it a point too obvious to require explanation, and I was not about to ask!) But I think the reason that helps in supporting the claim that what one's imagined speaker is saying is true is that it engages the general presumption that, where speakers are not basing their claims on some false beliefs they have about underlying matters of fact, how they naturally and appropriately describe a situation, especially by means of very common words, will be a true description.

if the subject asks us whether p is true, we will answer, 'Well, that's one of the things we're hoping to determine here.')

3. Truth/Falsity Asymmetry

Partly because they can be so reinforced by the naturalness and appropriateness of the claims involved, intuitions we have that claims are true are generally more trustworthy than are intuitions we might have to the effect that claims are false.[4] Of course, when a speaker makes a claim that is, and is from her own point of view, false, her claim will be improper as well as false. But this impropriety cannot be used to buttress the intuition that the claim is false with the same security that one can use the appropriateness of a claim to reinforce the intuition that the claim is true, because there is not nearly as strong a presumption that inappropriate claims are false as there is that appropriate claims are true. As David Lewis points out, 'There are ever so many reasons why it might be inappropriate to say something true. It might be irrelevant to the conversation, it might convey a false hint, it might be known already to all concerned...'[5] And though Lewis does not go on to explicitly say so, the comparative point he is making depends on, and he therefore implies that, it is not nearly as likely that an appropriately made claim will be false. And that implied point seems right. For it seems that, except where we engage in special (and fairly easily identifiable) practices of misdirection, like irony or hyperbole, we should seek to avoid asserting falsehoods, and we will thus be speaking improperly if we assert what is, from our own point of view, false. We should avoid falsehood not only in what we imply, but (especially!) in what we actually assert. So if a serious assertion is appropriate, that strongly supports the intuition that it is true, at least from the speaker's point of view. (And if the speaker isn't mistaken about any underlying facts, that the claim is simply true.)

Given that we really do (and with seeming truth and propriety) deny 'knowledge' in situations like HIGH, the contextualist is wise, then, to construct that case so that the speaker does there deny that the subject 'knows', rather than having the speaker ascribe knowledge in a seemingly false way.

---

[4] This of course is not to claim that intuitions to the effect that certain assertions are false are never worthy of credence. In fact, in section 15 of Chapter 1, for instance, I myself utilize a judgement that a claim to 'know' something is false as a premiss in an argument against contrastivism.

[5] Lewis (2000: 196); ellipsis in original. I advocated the asymmetry urged in this section in correspondence with Lewis in the early 1990s, where I argued that Lewis's account of the relation between 'might' and 'would' counterfactual conditionals was in particularly bad trouble because it ruled as false certain claims that intuitively seemed true and that were perfectly appropriate, while its key rival (which I advocated) merely had to rule true some claims that seemed wrong. This later came out in my (1994: at 417–18). Lewis might have already accepted that asymmetry at the time. As I recall, he raised no methodological objection to me in his correspondence.

Because of the comparative point made in the above paragraph, any intuition we might have that such a speaker's positive ascription of 'knowledge' is or would be false in HIGH would be more suspect than is the intuition that contextualists in fact rely on—that the denial of knowledge he actually puts in his speaker's mouth is true.

For this and some other reasons,[6] it greatly helps the contextualist's case that we do (actually, and with seeming propriety and truth) go so far as to deny 'knowledge' in situations like HIGH. Some of these other reasons will be presented in Chapter 3, below, especially in section 11, where we'll see that the above reason for favoring intuitions of truth over those of falsehood fits into a network of considerations that work together in favor of the methodology I am following to a contextualist conclusion. But hopefully one can already sense that the case for contextualism would be at least considerably weaker if the evident facts about HIGH were only that we are reluctant to positively ascribe 'knowledge' there and/or that if a misguided speaker were to positively ascribe 'knowledge' in such a case, the ascriptions would seem false and/or inappropriate. That the actual situation is such that, beyond the above, speakers go so far as to deny 'knowledge' in HIGH cases, and do so truthfully and appropriately, to all appearances, certainly seems to considerably strengthen the case for contextualism.

4. The Best Cases: Standards Appropriate to Practical Context

Contextualists can differ among themselves about what types of features of the context of utterance really do affect the truth-conditions of knowledge

---

[6] Stanley (2005: 123) expresses a worry about my proposal to put a higher premium on respecting intuitions of truth over those of falsehood. The reason Stanley gives for his suspicion is that my methodological claim 'makes an uncomfortable distinction between intuitions we have that assertions of negated propositions are true (which are, according to [DeRose] trustworthy) and intuitions that assertions of non-negated propositions are false'. But Stanley doesn't elaborate why he finds this discrimination uncomfortable, nor does he say what he thinks is wrong with the reason I give for thinking intuitions of falsehood are more suspect. In fairness, Stanley is only registering, and not seeking fully to justify, his suspicions.

A reason someone might have for being wary of basing arguments on claims about at-least-apparent denials of 'knowledge' is that these apparent denials might not really be denials of 'knowledge', but might instead be cases of what is known as 'metalinguistic negation'. However, I think we often, and in the case at hand, have strong reasons for rejecting the hypothesis that speakers are merely engaging in metalinguistic negation and for holding that they are instead involved in genuine negation: See especially n. 34 of Chapter 3.

attributions and to what extent they do so. Thus, there can be pairs of cases, differing only in 'disputed' features—features that some contextualists think affect the standards that comprise a truth-conditions of knowledge-attributing sentences, but that other contextualists think have no such effect. Such pairs of cases do not provide the best tests for deciding between contextualism and invariantism. The best case pairs will differ with respect to as many of the features that plausibly affect the epistemic standards, and especially those features which most clearly appear to affect epistemic standards, as possible. It's about such pairs of cases that the intuitions supporting contextualism will be strongest. And, given the content of their position, invariantists must resist the intuitions supporting contextualism even with respect to those strongest cases. By contrast, it's of course open to contextualists to hold that there is no difference in truth-conditions between the cases in some of the case pairs.

So, what makes for a good pair of cases? I'll here go into only a couple of important ingredients. But, first, it makes the relevant pro-contextualist intuitions stronger and more stable if the elevated epistemic standards in HIGH are tied to a pressing and very practical concern, and thus seem reasonable given the situation, and if the lower epistemic standards in LOW are also reasonable given the different practical situation involved there. And it helps (as we'll explore more fully in the following section) if all the parties to the conversation accept as reasonable both the elevated standards of the one case and the lower standards of the other. This is so of both of the case pairs we looked at in Chapter 1. In the Bank Cases, for instance, one character (myself, as it happens) claims to know that a bank will be open on Saturday morning in LOW. This belief is true, and, recall, is based on quite solid grounds: I was at the bank just two weeks ago on a Saturday, and found that it was open until noon on Saturdays. Given the low-grade practical concerns involved, almost any speaker in my situation would claim to know the bank will be open on Saturday. And, supposing 'nothing funny' is going on (there has not been a recent rash of banks cutting their Saturday hours in the area, that the bank will in fact be open on Saturday, etc.), almost all of us would judge such a claim to know to be true. But in HIGH, disaster, not just disappointment, would ensue if we waited until Saturday only to find we were too late. Given all this, my wife seems reasonable in not being satisfied with my grounds, and, after reminding me of how much is at stake, in raising, as she does, the possibility that the bank may have changed its hours in the last couple of weeks. This possibility seems well worth worrying about, given the high stakes we are dealing with. Here I seem quite reasonable in admitting to her that I 'don't know' that the bank will be open, and in endeavoring to 'make sure'. Almost everyone

will accept this as a reasonable admission, and it will seem true to almost everyone.[7]

We can encounter very tricky cases, about which contextualists can disagree among themselves, if the actual practical situation that a speaker faces is importantly different from how he perceives his situation to be. Thus, if there were a variant of the Bank Cases in which it is in fact vitally important that I be right about whether the bank is open on Saturdays, but in which I do not realize that the stakes are so high, perhaps because I have forgotten that I have written the large and very important check, or perhaps because my wife was the one who wrote the check without my being aware of it, so that I claim to 'know' the bank is open on relatively meager grounds, some contextualists may think the standards remain low and my claim to 'know' is true, while others may think the raised stakes, even when I'm unaware of them, raise the standards for knowledge that govern the truth-conditions of my use of 'know(s)' and its cognates, so that my claim is false. Such cases may be useful for deciding among different versions of contextualism, but are to be avoided in deciding between contextualism and invariantism in the first place. Better for that latter purpose, which is our purpose here, to make LOW involve stakes that are in fact low and are clearly seen to be low, and to have HIGH involve stakes that are in fact high and are clearly seen to be high.

To avoid confusion, I should here explicitly add that while I believe that the best case pairs for establishing contextualism involve a marked difference in the stakes involved, I myself do not believe that such a difference in stakes is necessary for a difference in the semantic standards for 'know(s)'. In fact, I think that speakers are free to use standards even wildly inappropriate to the practical situations they face—for instance, to use low standards where they face (and even realize that they face) an extremely high-stakes situation in which it would be much wiser for them to employ much higher standards. But when speakers do use inappropriate, and especially wildly inappropriate, standards, our intuitions about the truth-values of their claims become very insecure, partly, I think, because it is easy to confuse the fact that there is *something* very wrong with their utterance (as will be the case when they are employing standards wildly inappropriate to the situation they face) for their claim's being false. It is therefore far better to test contextualism by means of case pairs in which the standards employed in each case are appropriate to the

---

[7] Of course, we may begin to doubt the intuitions above when we consider them together, wondering whether the claim to know in the first case and the admission that I don't know in the second can really both be true. But when the cases are considered individually, without worry about the other case, the intuitions are quite strong, and, in any case, the linguistic behavior displayed in the cases quite clearly does accurately reflect how 'I know'/'I don't know' is in fact used.

practical situation of the case in question. And, again, about even such cases, where the intuitions are quite secure, the invariantist will have to rule that at least one of the relevant intuitions is wrong.

5. The Best Cases: Cases Involving No Dispute, No Reversals, and No Exceedingly High, 'Philosophical' Standards

Contextualist theories of knowledge attributions are typically developed with an eye toward providing some kind of contextualist response to the problem of philosophical skepticism. But it is important to first develop support for contextualism from ordinary, non-philosophical cases, because, among other reasons, the relevant intuitions are not very strong where the 'high-standards' case is one involving a philosophical skeptic, who denies 'knowledge' in ways that are not tightly tied to appropriate practical concerns, and whose denials seem to be based on standards that are *exceedingly* high. Many will find their intuitions in revolt against these philosophical denials of 'knowledge'.

So, contextualists have wisely appealed to cases which don't involve philosophical skepticism. It does help if something that functions as a 'skeptical hypothesis' is inserted into the conversation in HIGH—a hypothesis that is perhaps too far-fetched to be considered in LOW. But these hypotheses are still far more moderate than are the ethereal possibilities on which philosophical skeptics base their hyperbolic doubts (brains in vats, or whatnot), and do seem worth considering under the somewhat heightened scrutiny appropriate to HIGH. Thus, as we've seen, in my high-standards Bank Case, my wife raises the possibility that the bank has changed its hours in the last two weeks, and in his similar Airport Case, Stewart Cohen has one of his characters raise the possibility that the itinerary they are relying on to tell whether their flight stops in Chicago might contain a misprint (1999: 58).

But the best cases for contextualism do not just avoid the hyperbolically far-fetched possibilities and the astoundingly high epistemic standards of philosophical skeptics. The best cases avoid there being any dispute among the parties to the conversation over whether the proposition in question is or is not 'known', and avoid any of the speakers reversing what they say on the matter of what is or is not 'known'.

When such things are avoided, the intuitions supportive of contextualism are strongest. That's to be expected, since contextualists themselves can, and

some do, have worries about how to handle such cases. In fact, as we will see below in Chapter 4, with respect to at least some cases of dispute, my own (contextualist) position is that the speakers involved are not speaking truths. I have worked this out mainly with an eye toward applying it to disputes between *philosophical* skeptics and their commonsensical opponents, but the position I advocate—the 'gap view', on which *neither* the philosophical skeptic who insists that his opponent does not 'know' something *nor* his commonsensical opponent, who insists that she does indeed 'know' the thing in question, are making true claims[8]—can be easily applied to disputes involving more moderate, non-philosophical 'skeptics' and their opponents, and I would be inclined to apply that view to many such cases. So it would be a mistake to base contextualism on cases involving such disputes, claiming that both parties are speaking the truth—cases about which even card-carrying contextualists might have their doubts.

Some opponents of contextualism, apparently viewing contextualism as supported mainly by a certain resolution to the problem of philosophical skepticism that it makes possible—a resolution on which the apparently disputing parties are both speaking truthfully—have complained that such a resolution is always available for any dispute, and so the contextualist solution is far too general a problem-solver: We can always claim that a key term involved in a hard-fought dispute is context-sensitive, and thus claim the advantage for our 'contextualism' that it dissolves this apparent dispute. As Timothy Williamson sums up his version of this complaint: 'Contextualism supplies a perfectly general strategy for resolving any apparent disagreement whatsoever. Since some disagreements are genuine, we should not always follow that strategy.'[9] But this charge seems damaging only if contextualism is being supported exclusively, or at least mainly, by its ability to provide such a resolution to the dispute between philosophical skeptics and their opponents. Wise contextualists, however, as we have seen, seek to support their contextualism by appeals to what transpires in ordinary, non-philosophical uses of 'know(s)', the most important examples of which are cases that involve no dispute among the parties whatsoever, before *applying* their contextualism to the problem of philosophical skepticism. Other disputes involving other key terms can be just as well solved by a contextualist approach only if, among other things, similar evidence of context-sensitivity is available for those terms. Now, contextualists will typically also claim that the treatment of the problem of philosophical skepticism that we are able to supply provides further support for

---

[8] See section 11 of Chapter 4.
[9] Williamson (2005a: 694–5); see also Feldman (1999: 103 ff.) and Stanley (2005: 122–3).

their views. But the solutions there appealed to are *not* typically dissolutions of disputes between philosophical skeptics and their opponents on which both parties are speaking the truth and that are achieved by a one-size-fits-all generic contextualism. Indeed, contextualists typically don't even accept solutions which so dissolve the dispute: None of Stewart Cohen, David Lewis, or myself is inclined toward what I call a 'multiple scoreboards' view on which both the philosophical skeptic and her opponent are speaking truthfully when they argue, but instead opt for some 'single scoreboard' view on which the dispute is quite real indeed, and it's quite impossible for both disputants to be speaking truthfully (see Chapter 4). In fact, on the 'single scoreboards' view I opt for, and on both of the views that seem to be Cohen's finalists,[10] *neither* the skeptic nor her opponent is speaking truthfully as they dispute! We solve (or try to solve, depending on whom you talk to) various skeptical paradoxes that are generated by skeptical arguments not by making both the disputants in a fight turn out to be making truthful claims, but by explaining why we are led into a paradoxical set of intuitions. Our explanations are built upon various specific standards-changing mechanisms we claim govern knowledge attributions, and so our solutions are not instances of some general contextualist strategy available wherever there is some dispute that is hard to adjudicate.

Unfortunately for the invariantist, the cases where varying epistemic standards are displayed are not limited to murky cases where there is some dispute over what is 'known', or cases in which some speaker is led to reverse what she has said. It is a feature of how we use 'know(s)' that sometimes heightened practical concerns make it seem appropriate for speakers to deny that subjects 'know' things so long as those subjects don't meet unusually demanding epistemic standards, and speakers then do deny that the subjects involve 'know' propositions that the speakers would ordinarily say those subjects do 'know'. And sometimes—indeed, ordinarily—when they do this, they meet with no resistance at all from their conversational partners, who have also sensed that unusually stringent standards are appropriate in the context, and have adjusted their use of 'know(s)' to these high standards, themselves saying, or at least being ready to say, with apparent truth and appropriateness, that they and other subjects do not 'know' things if those subjects fall short of the extraordinarily high standards in question, even if they meet more ordinary standards. These best high-standards cases involve no element of arm-twisting, arguing, or convincing, but only *informing*—one speaker informing another that, according to the appropriately heightened standards that obviously govern

---

[10] See Chapter 4, near the end of section 11.

their conversation, some subject does not 'know' something. Paired off against such high-standards cases, it's also a feature of our use of 'know(s)' that often, where unusually heightened practical concerns are not involved, speakers behave in ways that show that they are in a low-standards context, informing each other that the same subjects described above, or similarly positioned subjects, *do* 'know' those same things that they are alleged not to 'know' in the high-standards case. And where there is no dispute among the parties to the conversation in either of the cases; and where none of the speakers in either case is made to reverse herself, saying that something is not 'known' that she had earlier claimed was 'known', or vice versa; and where the different sets of standards that seem to govern the two cases both seem appropriate, given the practical situations involved in each case; and where none of the claims about who does or doesn't 'know' what is based on false beliefs the speakers have about underlying matters of fact; the intuitions that the relevant claims in each of the cases (the ascription of knowledge in LOW and the denial of knowledge in HIGH) are true can become extremely powerful indeed. And these strong intuitions of truth are buttressed by the presumption that what is properly said is true. And even with respect to these strongest cases, invariantists must deny one of these intuitions involved.

6. Problems with First-Person Cases

Thus far, in looking for the best cases to use in support of contextualism, I have been concerned with ways the cases can be constructed to make the contextualist's key premises—that the affirmation of knowledge in LOW and the denial of knowledge in HIGH are both true—as secure as possible. But it is also important that the cases be constructed so that those premises support a contextualist conclusion as powerfully as possible.[11] For a couple of reasons, it helps achieve that second goal if the cases involve third-person attributions (and denials) of knowledge, preferably to subjects not involved in the conversation in which the attributions and denials are made, rather than first-person claims to know and first-person admissions that the speaker herself does not know. Thus, contextualists make a mistake when, as I have done, we focus our attention on cases involving first-person claims.

---

[11] In his presentation of some variations of the Bank Cases, Stanley inserts a 'really' into the denial of knowledge; he has the speaker say, 'I guess I don't really know that the bank will be open tomorrow' (2004: 121; 200; cf. 2005: 5). This is an alteration of the Bank Cases that hurts their effectiveness by making it needlessly problematic that the truth of the claims in question supports contextualism. This becomes problematic because the function of the 'really' is quite unclear.

To see one problem with first-person cases, consider my Bank Cases. In LOW, a speaker (myself, as it happens) claims that he himself knows that the bank is open on Saturdays, while in HIGH, the speaker admits that he does not know that the bank is open on Saturdays. For reasons we've been considering, the contextualist's premises that these two claims are each true seem quite secure. But do they really support contextualism? I believe so, but that certainly can be questioned. I think these cases support contextualism because the best explanation for why both claims are true, despite their surface opposition to one another, is that the semantic standards for knowledge differ between the cases. And that, in turn, is based on my sense that the subject—who is identical with the speaker in first-person cases like these—is in an equally strong epistemic position with respect to the bank being open on Saturdays in the two cases. What's always seemed to me the most natural explanation for why the speaker can truthfully claim 'knowledge' in the one case but truthfully deny it in the other, then, is that what semantically counts as 'knowledge' has changed, so that the subject's same epistemic position in the two cases meets the truth-conditions for 'knowledge' in LOW, but not in HIGH.

But recall the differences between the cases that, according to the contextualist, contribute to the difference in standards between the cases: In HIGH, the stakes are higher, making it appropriate to worry about possibilities of error that would usually be ignored; one such possibility is actually mentioned by the speaker's conversational partner, and is accepted in the conversation as a legitimate worry; the speaker then is thinking about, and taking seriously, this possibility of error as he makes his claim; etc. The invariantist may claim that it is far from clear that none of these differences affects how well positioned the subject really is with respect to the bank being open on Saturdays. Perhaps both the positive claim and the denial of knowledge are true, not because the semantic standards for 'knowledge' vary from one case to the next, but because, though the same standards govern both cases, the speaker meets those standards in the so-called 'low-standards' case, but fails to meet those same standards in the so-called 'high-standards' case.

It's worth quickly noting that contextualism would be validated by mixed explanations, according to which the differences between the two cases induce both a change in what semantically counts as 'knowledge' and the strength of the speaker/subject's epistemic position. To make good on the worry currently under consideration, then, the invariantist needs to hold out for the possibility that it is *entirely* a matter of the speaker/subject's difference in epistemic position between the two cases that explains how both the positive and negative claim can be true. Nevertheless, this worry considerably weakens

the contextualist's argument, so it is wise for the contextualist to sidestep this worry if possible.

And there is another problem with first-person cases that is worth considering—and that recent work by Jeremy Fantl and Matthew McGrath (2002), John Hawthorne (2004), and Jason Stanley (2005) on SSI makes more pressing. Perhaps different epistemic standards *do* comprise truth-conditions for knowledge attributions in the contextualist's two first-person cases, but contextualism is *still* not true, because these differing standards can instead be explained by SSI. According to SSI, which we looked at briefly in Chapter 1 (section 8), and will look at much more carefully in Chapters 6 and 7, the varying standards that comprise a truth-condition of 'I know that p' are sensitive to factors that attach to the speaker as the putative subject of knowledge, rather than as the speaker of the knowledge attribution. That is, according to SSI, these factors of the subject's context determine a single set of standards that govern when the subject himself, or any other speaker, including those not engaged in conversation with the subject, can truthfully say that the subject 'knows'. Thus, we do not get the result that contextualists insist on: that one speaker can truthfully say the subject 'knows', while another speaker, in a different and more demanding context, can say that the subject does 'not know', even though the two speakers are speaking of the same subject knowing (or not knowing) the same proposition at the same time. But where the subject faces a different situation, different standards govern whether she, or anybody else, can truthfully say that she 'knows'. And the backer of SSI can claim that just that is transpiring in the contextualist's first-person cases: different standards do govern the two cases, and they comprise truth-conditions, and not just warranted-assertability conditions of the knowledge-ascribing and -denying sentences, but these different truth-conditions work in the way SSI claims, rather than as the contextualist would have us believe.

7. Third-Person Cases

Fortunately, then, for the contextualist, we use third-person attributions of and denials of knowledge in describing subjects who are no party to our conversation in a way that demonstrates the same variation in standards that is displayed by our first-person talk of what we ourselves do and do not 'know', and, because of this, third-person pairs of cases are available that are about as powerful as are the best first-person cases. The two discussions involved in the office example would serve well here, with Thelma at the Tavern

comprising LOW and Louise with the Police being HIGH. These examples are as convincing as the first-person Bank Cases, with both the affirmation of knowledge in LOW and the denial of knowledge in HIGH seeming as clearly true as the analogous verdicts about the Bank Cases, and both speakers even more clearly seem to be using 'know(s)' appropriately, just as in the Bank Cases. And our third-person office examples also have the other features we've looked at here that make for the best contextualist cases.

But it might be instructive here to instead use third-person versions of the Bank Cases, to show how easily third-person cases can be derived from many first-person cases. So, when starting from first-person cases, take LOW, but make a topic of conversation whether some third party, not present at the conversation, 'knows' whether the bank will be open, and construct LOW so that one of the speakers, but not the other, knows that this absent subject has the same grounds for thinking the bank will be open on Saturday as the speaker has: the absent subject, too, was at the bank just two weeks ago on Saturday, and found it open, etc. You may suppose that this speaker has just called the absent subject and found this all out, and is now reporting the results of his phone enquiry to the rest of the discussants. Suppose that the far-away subject is herself at the time in a fairly ordinary situation, not some high-stakes discussion about any very important matter, and that she does not herself have anything important riding on whether the bank is open on Saturday. Now, even though that absent subject, like the speaker, hasn't made any of the special checks that would be needed to rule out such relatively far-fetched possibilities as that the bank has changed its hours in the last two weeks, it's clear that an English speaker, facing such a situation, who is saying that he himself 'knows' that the bank will be open on Saturday, will also say that this absent subject 'knows' that fact as well. And when the contextualist has the speaker claim that the absent subject 'knows', the intuition that that claim is true will be as powerful as it is in a first-person case, and the intuition will be buttressed by the related considerations that the speaker is speaking as we all in fact do, and with propriety, as much as in the case of the first-person case.

Next, similarly convert HIGH so that the topic of conversation is whether some third party not present at the conversation 'knows' whether the bank will be open. You can suppose again that the absent subject is herself at the time in a low-standards setting, and is unaware of the fact that, far away, our speakers are talking about her. But this time, of course, our two speakers are in a high-stakes, high-standards case, and are wondering whether various people, including our absent subject, might 'know' that the bank will be open. The speakers themselves have the relatively meager grounds we've discussed:

They were at the bank just two weeks ago on Saturday, and found it open. But because so much is at stake, they are worrying about the possibility that the bank has changed its hours in the last two weeks, and are consequently saying that they 'don't know' that the bank will be open. But they would have counted themselves as 'knowing' if, for instance, they had been to the bank just today, had asked about whether the bank would be open on Saturday, and had been assured by a knowledgeable bank employee, who herself is working on Saturday, that the bank is still open on Saturdays, and will definitely be open tomorrow. Such relatively strong grounds would be sufficient for our speakers to claim to 'know' that the bank will be open. Though our speakers don't at present have such relatively strong grounds, one of them is now wondering whether some friend of theirs, who may have been at the bank today and might have some interest in finding out for sure whether it will open tomorrow, might 'know' whether the bank will be open on Saturday, with the thought that if this absent friend does 'know', our speakers might then be able to use this 'knower' as an informant on the matter, and might themselves come to 'know' that the bank will be open. Give the absent subject in question the same relatively meager grounds for believing that the bank will be open on Saturday as in the low-standards case, and construct this new HIGH case so that, as with our new LOW case, our other speaker, unlike the speaker who is doing the wondering, knows what kind of meager epistemic position this absent subject is in at her far-away location, as well as what kind of conversational context that subject is in. It is clear that an English speaker who faces this situation, and who is saying that he himself does 'not know', will say that our absent subject, whom the speaker knows has no more evidence than the speaker himself has, and who is being considered as a potential informant, similarly 'does not know' that the bank will be open, and the contextualist's premiss that this high-standards third-person denial of knowledge is true will be about as secure as is his premiss concerning the high-standards first-person case.

In the third person case we have just considered, the speaker applies unusually high standards, appropriate to the speaker's own high-stakes context, to an absent subject who is neither in such a high-stakes context nor believed by the speaker to be in a high-stakes situation. This illustrates how speakers often talk, applying standards appropriate to their own conversational situations to a subject, even though, given the subject's own situation, those are not standards that would be appropriate for the subject to use, were she to say whether or not she 'knows'. Among the situations in which speakers do this are contexts like the one we've just considered in which the speakers are discussing those absent subjects as potential informants to the speakers.

But there are other conversational situations where quite different conversational purposes are in play, and in which speakers will apply to absent subjects standards that are appropriate to the practical contexts *of those subjects*. This often happens when, for example, the speakers are discussing practical decisions the absent subjects face: 'She should do that only if she knows that…' When discussing absent subjects in such cases, it is exceedingly odd, and perhaps even downright wrong, to apply to the subjects standards that are not appropriate to the subject's situation. We will discuss these different third-person cases in Chapter 7, but it's worth noting their existence now. While, as we will see, these different cases can be easily handled by contextualism, and thus present no real obstacle to contextualism, they are not the type of cases that are driving our present argument in favor of contextualism. While there clearly are conversational situations where it is appropriate to apply to subjects standards appropriate to the situations of those subjects, it is equally evident that there are also cases, like our third-person version of the high-standards Bank Case, where speakers apply to subjects standards appropriate to the speaker's own contexts, even though they are not the standards that would be appropriate for the subjects to apply to themselves. And it is such cases where the speaker's standards prevail that we are now utilizing in our argument, realizing that there are other situations in which we talk differently in this regard.

Since the contextualist premises we are now considering concern third-person cases, the worries we looked at in the previous section seem not to apply—or at the very least are not as worrisome. To press the first worry, the invariantist will have to claim that the absent party is in a worse epistemic position in our so-called 'high-standards' case as compared with her position in our 'low-standards' case. But we've put the absent party in the extremely similar situations in the two cases! As with the first-person cases, she has the same grounds for thinking the bank is open in our two new cases. And in our new third-person cases, she is not confronted with heightened practical concerns and a conversational partner pressing relatively far-fetched possibilities of error in HIGH—things which might plausibly be thought to hurt her epistemic position. In fact, our subject's situation is identical in our two third-person cases except for the difference that, far away from her, different conversations, which she is and will remain oblivious to, are taking place about her. So, to press the first worry, an invariantist must maintain that how strongly our subject is positioned with respect to the bank being open on Saturdays is substantially affected by such factors as what kind of far-away conversation is taking place about the subject. That's quite implausible on its face. Far better, it seems, to explain how the positive and negative claims in these two

cases can both be true by supposing that the semantic standards, rather than the subject's ability to meet a single set of standards, differs between these two cases.

With regard to our second worry, these third-person cases provide a powerful objection—to my thinking, a killer objection—to SSI. We will press this objection against SSI more fully in Chapter 7, where we will look at SSI's best hope of escaping the trouble it faces here. But, for now, just note that SSI incorrectly predicts that in our new, third-person version of HIGH, our speaker will apply the lower standards to our subject. For SSI holds that the situation *of the subject* sets the standards that govern all attributions and denials of knowledge to that subject, and our speaker is aware of the facts that make the subject's context a low-standards one.[12] Yet it is clear that in situations like the one our speaker faces in our third-person version of HIGH, speakers will apply their own, higher standards, to the subject: If the speaker is counting himself as not 'knowing', then, when he is speaking of the subject as a potential informant in the way we have been imagining, and this subject is no better positioned than is the speaker, the speaker will say that the subject, too, 'does not know'. Thus, against very strong appearances, and against the methodology of ordinary language that we have looked at here, the defender of SSI will have to deny that the denial of knowledge in HIGH is true—or that the positive ascription of knowledge in LOW is true.

But what of the different type of third-person case, where speakers apply to subjects the standards appropriate to the practical situation of the subjects? Of course, SSI handles these different cases very nicely. But as we'll see in more detail in Chapter 7, contextualism can handle these different cases, too. Yet, of the two views, only contextualism can handle the cases we've been focusing on in this chapter, where, as clearly often happens, speakers apply to subjects standards appropriate to the speaker's own context, even though they would not be the standards appropriate for subjects to use to describe themselves. For only contextualism provides the flexibility to use either type of standard, while SSI demands that, no matter what their conversational purposes are, speakers

---

[12] It's a bit artificial to suppose that our speaker has as much information as I've given him about the subject's conversational situation. I did so to make the argument simpler. In more realistic situations, where the speaker doesn't know what standards govern the subject's context, SSI makes somewhat more complicated, but equally false, predictions: It predicts that the speaker will be agnostic about whether the absent subject 'knows', since he doesn't know what standards are appropriate to the subject's context. (In fact, since the standards the speaker applies are extraordinarily high, he should be highly dubious that these are the standards that govern whatever context the subject is in.) But that's not how speakers in fact behave in our speaker's situation. If the speaker knows that the subject doesn't meet the unusually high standards that govern the speaker's context, the speaker will say that the subject, like the speaker himself, 'does not know'.

must apply to subjects the standards appropriate to the subjects' situations. It's easy to see how our language could have operated so that SSI would have been true of it: Speakers would always apply to subjects the standards appropriate to the subjects' situations. In other words, there would be only one of the two types of third-person cases we've distinguished here. But as cases like the one we've been focusing on here show, our language in fact does not operate in this way.

8. The Importance of Arguments from Ordinary Language

Arguments from ordinary usage—appeals to what speakers of a language do or would say in applying the terms in question to particular situations (both positive and negative claims involving the term), appeals to which simple applications are or would be proper or improper for them to make, and appeals to intuitions as to the truth-values of those claims in particular situations—don't exhaust the considerations that need to be attended to in deciding among semantic theories, but do provide absolutely essential considerations.

   To introduce considerations I will return to in later chapters, this importance is most forcefully impressed upon me when I consider some obviously false theories about the truth-conditions of various sentences, and ask myself how (on what basis) *we know* that they are false, and the closely related, but perhaps even more important, question of what *makes* these theories false (in what does their falsehood consist)? So, to take one example of a theory that is obviously too demanding, and one that is obviously not demanding enough, consider, for instance, the crazed theory according to which a necessary condition for the truth of 'S is a physician' is that S be able to cure any conceivable illness instantaneously,[13] and the theory about the meaning of 'bachelor' like the traditional account, except that it omits the condition that S must be unmarried for 'S is a bachelor' to be true, insisting that married men can be truthfully (if perhaps for some reason improperly) called 'bachelors'. There are a couple of real clunkers! But in virtue of what is our language in fact such that these strange theories are not true of it? I'm of course not in a position to give a complete answer to this question, but, with respect to the strange theory

---

[13] This is based on an example used by Barry Stroud (Stroud 1984: 40), who in turn borrowed it from Paul Edwards.

concerning 'physician', it seems eminently reasonable to suppose that such facts as these, regarding our use, in thought and speech, of the term 'physician', are centrally involved: that we take to be physicians many licensed practitioners of medicine who don't satisfy the demanding requirement alleged; that we seriously describe these people as being 'physicians'; that we don't deny that these people are 'physicians'; that claims to the effect that these people are 'physicians' intuitively strike us as true; etc. It's no doubt largely in virtue of such facts as these that the traditional view, rather than the bizarre conjecture we are considering, is true of our language: The correctness of the traditional view largely consists in such facts. And these facts also seem to provide us with our best reasons or evidence for accepting the traditional, rather than the strange, hypothesis regarding the semantics of 'physician'. Analogous points hold for why a traditional theory of 'bachelor', rather than the bizarre theory we are considering, is true, and for how we know it's true.

Moving to context-sensitivity, such facts about ordinary usage also provide us with our primary, most important, and best evidence that clearly context-sensitive terms like 'tall' are context-sensitive in the way that we take them to be, and are also that in which the context-sensitivity of those terms consist. Ignoring many subtleties, 'tall' is context-sensitive in the way we take it to be largely because speakers do (in fact, with propriety, and with apparent truth) seriously say that things are 'tall' in 'low-standards' contexts when those items meet certain moderate standards for height, even if they don't meet still higher standards, but will, in 'high-standards' contexts, seriously deny that those same things are 'tall', reserving the application of that adjective only for items that meet the more demanding standard for height. In such facts about ordinary usage we have some of the best possible type of evidence you could ask for that 'tall' is context-sensitive.

But likewise then for 'know(s)'. Data of the type we've been considering here provide us with the best possible type of evidence you could ask for that 'know(s)' is context-sensitive in at least roughly the way contextualists claim it is. 'Know(s)' is context-sensitive in that way largely because speakers in some contexts do (in fact, with propriety, and with apparent truth) seriously describe subjects as 'knowing' propositions when those subjects meet certain moderate epistemic standards with respect to the propositions in question, even if they don't meet still higher epistemic standards, but will, in other contexts, go so far as to (in fact, with propriety, and with apparent truth) seriously deny that such subjects 'know' such things, reserving the ascription of 'knowledge' only for subjects that meet some more demanding epistemic standard.

Note that this argument from ordinary usage for contextualism is not an indirect argument that takes as its premiss that some other term, like 'tall',

is context-sensitive, and then argues that because 'know(s)' is so similar to 'tall', 'know(s)' too is context-sensitive. Such an indirect argument would be very insecure, in my opinion, because, while there are very important similarities between the behaviors of 'tall' and 'know(s)', there are also many important differences. (We are, after all, comparing a verb with an adjective!) Rather, this case for contextualism is *direct* in that it adduces the same kinds of considerations for the context-sensitivity of 'know(s)' as are available for the context-sensitivity of the likes of 'tall'. Now, I will want to ask those who resist the argument on what basis they think we can know that 'tall' is context-sensitive, and why they think we don't have the same types of grounds for thinking 'know(s)' is context-sensitive, but this is not to revert to the indirect argument. Rather, it is an effort to compare the direct argument for contextualism with the direct grounds available for the context-sensitivity of 'tall', and in so doing to give a higher-level argument that the direct argument for contextualism is solid.

As for views like SSI, relative to the types of considerations we have here considered, it is easy to imagine what our language would have been like so that SSI would have been true of it. In short, and crudely skipping lots of details, we would always apply to subjects standards appropriate to those subjects' contexts (and would be agnostic about whether those subjects 'know' where we were problematically unsure about the nature of their practical and conversational, as well as epistemic, situation). That we do not talk about subjects in that way is, in large part, what makes SSI false—it is that in which the falsehood of SSI consists—and also provides us with our best—and quite decisive, in my view—evidence that SSI is false. But I will make this case in more detail in Chapter 7.

The positive, direct argument for contextualism we have explored here utilizes fairly straightforward considerations about how simple claims about what subjects do and do not 'know' are in fact, and with seeming propriety and truth, used in ordinary talk. There are more complex matters to consider, some of which have been used in arguments against contextualism—matters like how knowledge attributions and denials behave when embedded in speech reports and propositional attitude reports, what metalinguistic claims speakers will make, and how we tend to judge how the content of one claim compares with that of another. These matters will be investigated below in Chapter 5, where I argue that the arguments against contextualism based on such considerations fail. In fact, I will suggest that, if anything, investigation of the more complex features of the behavior of 'know(s)' tends to confirm, rather than refute, contextualism, at least so far as anyone I know of has shown.

However, I do not base contextualism on any argument from such more complex considerations. My main positive argument for contextualism is the one we have begun to investigate here in this chapter, based on relatively straightforward data concerning simple positive and negative claims involving 'know(s)'. The main resistance to this argument has been the 'warranted assertability objection' that will be addressed in Chapter 3, below, where we also look at another powerful positive argument in favor of contextualism.

Appendix: Similar Arguments for Other Contextualisms: But I Still Don't Know Who Hong Oak Yun Is!

It is worth briefly considering whether we have grounds similar to those I am endorsing here for other forms of contextualism. This may help clarify my argumentative strategy, and may give some indication of how much contextualism one is likely to be ending up with if one accepts my reasoning for the contextualism that is the concern of this volume.

The epistemologist David Braun opens his (2006) with these provocative words:

Hong Oak Yun is a person who is over three inches tall. And now you know who Hong Oak Yun is.   (2006: 24)

Braun bases this, first, on the following analysis:

> *Knowing Q*: If Q is a semantic question, then *X* knows Q iff *X* knows a proposition that answers Q. (2006: 32)

On this analysis, for example, you know who Larry kicked if and only if you know to be true some proposition that answers the question, 'Who did Larry kick?'; you know where Larry lives if and only if you know to be true some proposition that answers the question, 'Where does Larry live?'; and you know who Hong Oak Yun is if and only if you know to be true some proposition that answers the question, 'Who is Hong Oak Yun?' But when does a proposition answer a semantic question? The second basis for Braun's position is his 'Information Provision analysis' (IP analysis) of that relation:

> A proposition answers a semantic question iff it provides information about the question's subject matter. (2006: 27)

Now, I still have some unresolved questions about just when it is, for Braun, that a proposition provides information about a question's subject matter—various examples of minimal-at-best 'aboutness' about which I'm not sure how Braun would rule—but seeing how Braun applies this criterion I think has given me enough of an idea of what he has in mind to proceed. *Larry kicked Mary* does provide information about the subject

matter of the question 'Who did Larry kick?', and so it answers that question. But notice that Braun's IP analysis doesn't require that a proposition provide particularly useful or substantial information for it to answer a question; 'Who did Larry kick?' is also answered by *Larry kicked someone who is over three inches tall*, on this analysis. Braun applies this analysis to his focus case of knowing who:

(9a) Who is Hong Oak Yun?

(9b) Hong Oak Yun is a person who is over three inches tall.

It is very unlikely that any actual enquirer who posed the question expressed by (9a) would be satisfied with the proposition expressed by (9b), and only a smart aleck would respond with an assertion of the proposition expressed by (9b). Nevertheless, that proposition provides genuine information about the subject matter of (9a), so it is an answer according to the IP analysis.   (2006: 30)

And, thus, by Braun's analysis of *knowing Q*, anyone who knows that (9b) is true (as Braun's readers are supposed to because he has just told them) knows who Hong Oak Yun is. Indeed, to anticipate a handy piece of technical terminology we will use in Chapter 7, on Braun's analysis, anyone who knows that (9b) is true *simply knows* who Hong Oak Yun is, by which we mean that any speaker, whatever their context, who says that such a person 'knows who Hong Oak Yun is', will be asserting a truth, and any speaker will be asserting a falsehood who says that such a person 'does not know who Hong Oak Yun is'.

I present Braun's account because he goes on to evaluate his very liberal, invariantist account of knowing Q against a rival contextualist account. It will be instructive to compare Braun's construal of this contest with my own, since his view of the various considerations in play illustrates aspects of many philosophers' approach to questions over context-sensitivity.

Consider a contextualist account that agrees with Braun that knowing who someone is is a matter of knowing to be true propositions that provide information relevant to the question of who that person is. But instead of agreeing with Braun that S's knowing to be true *any* such proposition is always sufficient for the truth of 'S knows who X is', the contextualist holds that what propositions S must know to make such a claim true is a matter that varies from context to context: In some contexts, for instance, S's knowing that George W. Bush is a famous American is enough to make true an assertion of 'S knows who George W. Bush is', but in other contexts, 'S knows who George W. Bush is' is false if the only substantial knowledge that S has about Bush is that he is a famous American.[14] And even if it is allowed there may be some bizarre

---

[14] Of course, there are importantly different ways of filling out such a contextualist account of the semantics of ascriptions of 'knowing who'. And many other issues come up for different sentences that fall under the very broad category of 'knows Q' sentences. We are here looking only at a bare-bones theory of 'knowing who' designed to make the most direct contact with Braun's invariantist account. For a good look at many of the relevant issues that arise in connection with 'knowing who', see Boër and Lycan (1986).

contexts in which S's knowing that Hong Oak Yun is a person who is over three inches tall suffices for the truth of an assertion of 'S knows who Hong Oak Yun is', there certainly are also many more normal contexts in which such meager knowledge of Yun's height is not enough to make such a claim true.

Braun says—correctly, in my view—that the above described contextualism concerning 'knowing who' 'strongly parallels' the contextualism we are investigating in this volume (which we will for now call 'standards contextualism'), with many of the arguments and counter-arguments for and against one contextualism largely paralleling those of the other (2006: 35). And the reader will probably be not at all surprised to learn that I accept this contextualism concerning 'knowing who' (and, more generally, I tend toward contextualist accounts about knowing Q, though various of these forms of knowledge raise important issues that are particular to themselves), as well as the standards contextualism I'm defending in this volume. But *why* accept such a contextualist account? Braun writes:

The primary evidence in favor of Contextualism is the fact that our judgments about the truth-value of sentences of the form '*A* knows *Q*' vary from context to context. In some contexts, we might think that 'Karen knows who Mark Twain is' is true because she knows that Mark Twain is an author. In other contexts, we might judge that it is not true because she does not know that Mark Twain wrote *Huckleberry Finn*. Contextualists hold that these sentences genuinely do vary in truth-value (and semantic content) from context to context. (2006: 34)

Having just encountered my description in this chapter of the case for standards contextualism, the reader should notice the discrepancy between Braun's description of the 'primary evidence' for contextualism and the case for standards contextualism as I am defending it. Braun sees the contextualist as arguing primarily from different truth-values for positive assertions of 'S knows who X is' in different contexts. So, Braun sees contextualists as arguing from pairs of cases (or perhaps larger sets of cases) in which, for instance, a subject who has all the same knowledge about Mark Twain in the two cases is described by speakers in quite different contexts as 'knowing who Mark Twain is', but in which that positive ascription seems true in one of the contexts, but false in the other. And Braun sees the contextualist as arguing from the premises that the one ascription is true and the other is false to a contextualist conclusion.

However faithful Braun's presentation of the case may or may not be to the contextualists whom he is addressing,[15] I would judge such a case for contextualism to be quite weak, due to the shakiness of the premiss that the ascription in one of the contexts is false. First of all, to cite a consideration that I haven't yet raised, it may be more difficult than one might think to devise a case that is otherwise suitable to the argument and in which it really does strike us as intuitively clear that the positive

---

[15] The contextualists Braun cites are Boër and Lycan (1975, 1986), Quine (1977), Kvart (1982), and Salmon (1987).

ascription in that case is false, because there is pressure on us as interpreters of the ascription to understand it as having a content that makes it true, due to the operation of what David Lewis calls a 'rule of accommodation' (Lewis 1979: esp. 346–7). And whatever intuition that we might nonetheless have that the ascription is false would be a relatively suspect one, for reasons we have already seen in sections 2 and 3 of this chapter.

Fortunately for contextualism concerning 'knowing who', the case for the view as Braun conceives it is not only quite weak, but is also needlessly weak: the contextualist has a much stronger argument available to him. Instead of pairing a case in which, for example, a first speaker seems to truthfully describe a subject as 'knowing who Mark Twain is' with a case in which a second speaker seems to falsely assert the same positive sentence about the same subject (or about a similarly positioned subject), it will be easy for the contextualist to instead construct a case in which a second speaker seems to truthfully and quite appropriately *deny* such knowledge to the same subject that the first speaker positively describes (or a similarly positioned subject), saying that the subject 'does not know who Mark Twain is'. If the case is constructed well, the intuition that this denial is true should strike us as quite clear, as opposed to the more cloudy status of whatever intuition we might have concerning the premiss used by Braun's contextualist. And, for the reasons we considered in sections 2 and 3 of this chapter, this intuition is reinforced by the fact that speakers do, with apparent propriety (as well as apparent truth), assert these denials of knowledge of who someone is in the relevant cases, even where these assertions are not based on any false beliefs the speakers have about underlying matters of fact.

A defender of Braun might complain with some plausibility that I'm reading the above indented quotation uncharitably: Perhaps, to go beyond the letter of what he wrote, Braun meant to include *denials* of knowledge of who someone is among the claims the truth-values of which his contextualist is appealing to as 'primary evidence'. (I would not be at all surprised to learn that at various points in this very volume, I write simply of sentences or of claims 'of the form "S knows that p"', where I mean to indicate both sentences or claims of that positive form and also those of the form 'S does not know that p'.)

But, at any rate, I think it is safe to say that Braun is not sensitive to how much the contextualist's case is strengthened when the contextualist shifts his basis from how it is presented by the letter of Braun's description to the case I advocate. Braun claims to neutralize 'the primary evidence in favor of contextualism' by providing an alternative account of our relevant judgements that he thinks is just as good as the contextualist's account:

As I said, the primary evidence in favor of contextualism is the variability of our judgments about the truth-value of sentences [of] the form '*A* knows that *Q*', for instance, '*A* knows who *N* is'. The contextualist explains these judgments by supposing that the contents and truth-values of such sentences vary from context to context. I have already (in effect) provided an alternative explanation. The semantic content of the attribution does not vary from context to context. But a cooperative speaker *S* tends to say and think, at time *t*

(where t is the time of S's context), that X knows who Y is only if X knows an answer to the question *of the sort that will satisfy an inquirer who is salient to S at t*, for instance, S herself, or her auditor. Since salient inquirers vary from context to context, S may (mistakenly) say and think that X knows who Y is in one context, but say and think the negation of this same proposition in another.

Both theories explain contextual variability of judgments. Both attribute mistakes to speakers. Contextualism attributes ignorance of context-sensitivity. My theory attributes a (sometimes) failure to distinguish between responses that are answers and responses that (also) satisfy the inquirer.

The evidence for Contextualism does not favor it over the non-contextual, non-interest-relative theory I advocate.   (2006: 35)

Braun does indicate at the end of the first paragraph of the above passage that he hopes to be explaining speakers' denials of knowledge of who someone is, as well as positive ascriptions of such knowledge. However, the principle he explicitly appeals to—the sentence that begins with 'But a cooperative speaker…'—states only a necessary condition for positive ascriptions of such knowledge, and so is silent on denials. But perhaps I am being uncharitable again, and we are clearly supposed to understand the principle as if it had something like '…and otherwise, speakers tend to say and think that X does *not* know who Y is' appended to its end?

In either case (but especially if the principle is to be understood as having the addition about denials appended to it), Braun faces this question that seems to me exceedingly pressing: If not here, where *should* we take seriously our intuitions about the truth-values of assertions, as opposed to giving up on them whenever some perhaps very counter-intuitive theory comes along that can 'succeed' in explaining how we use a term, so long as we don't count against the success of the rival's explanation that it implicates speakers in widespread and systematic falsehood in their use of the term and issues very counter-intuitive results as to those truth-values? In other words: Just when *does* it count seriously against the success of a theory's account of our use of the term that involves us in such widespread and systematic falsehood and issues such counter-intuitive results? Is our evidence from variable usage for a contextualist account of 'tall' neutralized by the fact that it is easy to construct an invariantist theory that can account for our variable usage of 'tall' that would seem to be about as successful as Braun's invariantist theory about knowing who someone is? Such a theory would insist that 'tall' is always semantically governed by the same standard, but would account for our variable usage of the term via a claim that speakers tend to say and think that X is tall iff X meets some standard for height that is salient to them as they make their judgement or assertion, and will otherwise tend to deny that X is tall? Or what about a quite strange invariantist account of 'here' (which, I admit, would be less plausible than Braun's theory), according to which, so far as semantic content goes, that term always refers to the same place, but explains our variable usage by means of a claim that speakers tend to say and think that X is here iff X is at the speaker's location, and they will otherwise tend to deny that X is here? Is that theory sufficiently damaged by its counter-intuitive rulings about the truth-values of our uses of 'here' and by the

widespread and systematic falsehood it charges speakers with in that usage, or must we really hunt around for other grounds on which to reject it?

In the second paragraph of the above quotation, Braun does seem to indicate that it's a mark against his theory that it attributes mistakes to speakers, but it isn't clear how serious a problem he takes this to be, because he quickly avers that this problem is balanced by a different mistake that contextualism must charge speakers with. Since he then immediately concludes (in the third paragraph of the above) that the evidence under consideration does not even favor contextualism over his own view, one is led to conclude that Braun thinks the mistake contextualism saddles speakers with is at least about as serious as are the mistakes he attributes to speakers. But is that because both theories attribute extremely serious mistakes, and thereby have extremely serious marks against them, or is this a balance of not-so-serious side considerations? It's hard to tell.

But, at any rate, it is also hard to see how it can be that contextualism is in anything close to the trouble that Braun's theory is in here. Braun's theory issues wildly counter-intuitive results about the truth-values of a very broad range of our uses of the relevant sentences. (Other possible invariantist accounts would score better here, but all would compare unfavorably with contextualism on this matter.) What balances that? Braun assures us that both theories 'attribute mistakes to speakers', but all we get in the case of contextualism is this: 'Contextualism attributes ignorance of context-sensitivity.' Now, first, ignorance seems not yet to be a *mistake*. Braun's theory attributes to speakers all manner of wildly false beliefs and assertions about whether various people know who various people are. What false beliefs or assertions does contextualism implicate us in? Braun seems not to be able to bring himself to claim that ordinary speakers actually believe that the relevant claims are context-insensitive in the relevant sense—for the very good reason, I think, that most speakers have not considered the matter and have no such belief. So he charges contextualism with charging speakers of *ignorance* of context-sensitivity. But, to anticipate how I will respond in Chapter 5 (see sections 3 and 8 of that chapter) to a similar objection aimed at standards contextualism, it is very far from clear that Braun's invariantism has any advantage at all over contextualism here, much less an advantage capable of achieving the balance he seeks. Here, I take it, are the relevant facts. First, as I've already noted, most speakers have not considered whether the claims in question are context-sensitive in the relevant way, and have no opinion on the matter, one way or the other. Second, most speakers (in my experience, the clear majority of them) who *have* considered this question and formed an opinion about it believe that the claims *are* context-sensitive; these are mostly philosophers and linguists, in my experience. Third, most ordinary speakers (and the strong majority of them, in my admittedly very limited experience of trying this out) who have the issue explained to them and who come to have an opinion about it come to think that the relevant claims *are* context-sensitive.[16] In light of these facts, it is far from

---

[16] The second and third facts I cite describe a situation that is considerably more favorable to contextualism about knowing who someone is than is the analogous situation concerning standards

clear, to put matters mildly, that contextualism implicates speakers in an ignorance of context-sensitivity that is any more worrisome than is the apparently somewhat more robust (in that it's an ignorance that seems more often to survive its bearer's coming to understand and consider the matter) ignorance of context-*in*sensitivity that Braun must charge them with, much less that he has an advantage here that's anywhere close to being strong enough to achieve the balance he needs. He will have a hard enough time countering a plausible charge that his theory faces yet another disadvantage here, I believe.

This quick look at another form of contextualism, I hope, gives some readers who need it a better sense of the potential for the general applicability and importance of the lines of argument we are considering in this volume. And I hope that my brief wrestling with Braun here will make more vivid the importance of the task of determining which judgements about the truth-values of the ordinary assertions speakers are inclined to make should be taken seriously and which assertions we should strive to make come out true as we construct and evaluate semantic theories. When should we consider a serious mark against a theory that it fails on these fronts? I have said that this is a very pressing question that faces Braun, but it faces us as well. I have begun that task in the current chapter, and will continue it in the following chapter, especially in sections 3 and 11 of Chapter 3, arguing that the premises used in support of standards contextualism are of the type that should be taken very seriously indeed. (And the same conclusion would hold for the premises used in the best arguments for the contextualism that Braun opposes.) Braun's essay also points forward to what's to follow in this volume because he raises against contextualism concerning knowing Q objections that, as he realizes, are analogous to objections that have been raised against standards contextualism, and that I will rebut in what follows (especially in Chapter 5).[17]

---

contextualism that I address in Chapter 5. Thus, in Chapter 5, I do not make such strong claims in defense of standards contextualism—though I think the facts about speaker's opinions about whether context-sensitivity obtains are still favorable enough to contextualism to provide a sufficiently strong response to this objection to contextualism. For this reason, I suspect that the contextualism concerning knowing who someone is that Braun opposes is even more defensible than is the standards contextualism that I am defending in this volume. Another respect in which contextualism regarding knowing who someone is may at least appear to be on more stable ground than is the standards contextualism that I am here defending is that the case from ordinary usage for the latter, but not for the former, contextualism is subject to the 'generality objection' that we will encounter in Chapter 3, section 4—though I ultimately think this objection can be defeated, too (see Chapter 3, section 10). On the other side of the ledger, however, standards contextualism can also be supported by an additional argument that I think is quite powerful and that we will look at in Chapter 3 (especially section 8), while no argument like that seems available for contextualism concerning knowing who someone is.

[17] At (2006: 37–8), Braun presses an objection that he says is 'peculiar to Contextualism about "knows who"'. I think that this would actually be an objection that would apply more generally to contextualism about knowing Q, and I suspect that's what Braun meant; what's important about it's being 'peculiar' is that it doesn't apply to the standards contextualism that I am defending here. At any rate, since this objection isn't aimed at standards contextualism, it is not among the objections I respond to in the chapters to follow. Even so, I think that some of my maneuvers in Chapter 7 will give most

It may help some readers for me to even more quickly discuss another form of contextualism that is very prominent these days: contextualist accounts (or what Stephen Schiffer calls 'hidden indexical' accounts (1992)) of belief attributions. Briefly and roughly, according to such theories, an assertion of 'S believes that p' is true iff S believes that p under a mode of presentation that is contextually relevant to the assertion. Since which modes of presentation are contextually relevant can vary from context to context, this theory rules that it can happen that one speaker truthfully asserts that 'S believes that p', while another speaker in another context, but speaking of the same S and the same p at the same time, truthfully asserts that 'S does not believe that p'. I accept such a theory, and I do so precisely because the type of grounds I am advocating in the case of standards contextualism also seem to be available, and seem to me to be convincing, in the case of these theories of belief attributions. The grounds I have in mind concern the seeming appropriateness and truth of both positive ascriptions of belief and denials of such belief, made by different speakers, both of whom are describing the same subject's relation to the same proposition, and both of whom have a correct view of all the facts relevant to whether the subject in question has the belief in question, but who are in different conversational contexts and have different conversational interests in the matter of whether S believes that p.

I should quickly add that, though many of the considerations that are important to evaluations of these different forms of contextualism seem to be very similar, there are also important features of the relevant arguments that will be particular to one or the other contextualism, and even where the considerations seem to closely parallel each other, they may more forcefully support a contextualist or anti-contextualist conclusion in the argument over one form of contextualism than in an argument over another. Thus, for instance, in n. 16, I point out some important differences between the arguments for standards contextualism and for contextualism concerning 'knowing who', and there seem to me to be even more important differences when it comes to the issue of modes of presentation contextualism concerning belief attributions. I have not here attempted a serious evaluation of these other contextualisms; we have barely scratched the surface of these issues. But I hope that this brief glance at them is enough to indicate the potential for at least many of the lines of argument we are considering in this volume to be of more general importance.

And, finally, for those who may be beginning to suspect that I am very quick to accept theses according to which the semantic content of sentences varies with context, it may be useful to some readers to briefly mention a well-known issue where the kind of argument I'm advocating is not available, and in which I don't endorse the conclusion that the content of sentences of a certain type varies with context. In some well-known papers (1966, 1968, 1978), Keith Donnellan draws a distinction between 'attributive' and 'referential' functions or uses of definite descriptions, which Donnellan handily summarizes as follows:

readers a pretty good idea of how I would respond on behalf of contextualism concerning knowing Q to this objection of Braun's.

A speaker who uses a definite description attributively in an assertion states something about whoever or whatever is the so-and-so. A speaker who uses a definite description referentially in an assertion, on the other hand, uses the description to enable his audience to pick out whom or what he is talking about and states something about that person or thing. In the first case the definite description might be said to occur essentially, for the speaker wishes to assert something about whatever or whoever fits that description; but in the referential uses the definite description is merely one tool for doing a certain job—calling attention to a person or thing—and in general any other device for doing the same job, another description or a name, would do as well.   (Donnellan 1966: 285)

Though Donnellan himself is a bit cagey about the matter, his work suggests, at least to some, that there is a semantic ambiguity here—that the content of what is asserted by a sentence that contains a definite description can vary depending on which of these two uses of definite descriptions the sentence contains.[18] As I suggest at the start of section 12 of Chapter 1, theses according to which a term has two senses—like the 'two senses of "knows"' theories discussed there and the claim that definite descriptions are ambiguous currently under discussion—can be viewed as limiting cases of contextualism,[19] and, in any case, such cases share with contextualism the feature that what is asserted by the use of a sentence varies from context to context.

   But in this case, we don't get the kind of evidence for this variation in content that I am looking for. The problem is with the cases where Donnellan would classify the use of definite descriptions as referential. Here, we don't get the kind of evidence we need that the truth-value of the assertion differs from what it would be if the definite description were used attributively. In his well-known discussion of Donnellan, Saul Kripke discusses the following case:

Someone sees a woman with a man. Taking the man to be her husband, and observing his attitude towards her, he says, 'Her husband is kind to her,' and someone else may nod,

---

[18] In his well-known discussion of Donnellan (Kripke 1977), Saul Kripke complains about this caginess. In a later paper, Donnellan writes: 'It might be thought, however, that if the position of this paper were correct that an ambiguity in the definite article would at least be suggested and that it is intuitively very implausible to suppose such an ambiguity' (Donnellan 1978: 41), and he goes on to defend himself from this charge. Note, however, that Donnellan does not even say in his own voice that his position so much as suggests such an ambiguity, much less that he posits such an ambiguity.

[19] Here it may be very important that the two ambiguity theories we are considering each posit two senses of the relevant terms that, though they are distinct, are quite closely related to one another. The alleged 'strong' and 'weak' senses of 'know(s)', in particular, are very closely related: both require true belief of the relevant proposition in question, and even where the requirements of the two alleged senses differ, the strong sense can be viewed as requiring more of the same kind of thing that is required by the weak sense. The paradigms of ambiguity, at least for many, are cases like 'bank', which is ambiguous between designating a certain kind of financial institution, and a river bank. I fear that when such people hear that a theory posits an ambiguity, their operative thoughts run something like: 'Oh, according to this theory, this term is like "bank"/"bank"', when in fact there may be very important differences in the behavior of terms that have two closely related senses as compared with cases of ambiguity like 'bank'. It's when a theory posits two closely related senses of a term that it seems most appropriate to think of it as a limiting case of contextualism.

'Yes, he seems to be.' Suppose the man in question is not her husband. Suppose he is her lover, to whom she has been driven precisely by her husband's cruelty.   (Kripke 1977: 7)

Supposing the lover, but not the husband, is kind to the woman, what we need from such a case to support the claim that definite descriptions have a separate referential sense is a premiss to the effect that the speaker's claim in this case is true. But Kripke writes:

If Donnellan had roundly asserted that the quoted statement is true if and only if the *lover* is kind to her, regardless of the kindness of the husband, the issue between him and Russell would be clearly joined. But Donnellan doesn't say this: rather he says that the speaker has referred to a certain person, the lover, and said *of him* that he is kind to her. But if we ask, 'Yes, but was the statement he made true?', Donnellan would hedge. …

Donnellan's cautious refusal to say, under the circumstances mentioned, that 'Her husband is kind to her' is true, seems nevertheless to be intuitively correct. The man to whom the speaker refers is—let us suppose—kind to her. But it seems hard for us to say that when he uttered, 'Her husband is kind to her,' it expressed a truth, if *we* believe that her husband is unkind to her.   (Kripke 1977: 12, 13)

Comparing the evidence for a referential sense of definite descriptions with the kind of evidence I seek (and think I find in the case of standards contextualism), then, the obvious and simple shortcoming of the former is that we just don't have the intuitions about the truth-values of the relevant claims in the cases that we need: I certainly agree with Kripke that it is far from intuitively clear that the speaker in this case asserts a truth—and the intuitive situation seems similar in other cases that Donnellan would classify as referential.[20]

---

[20] I'd also add that it's far from intuitively clear that the assertion is false. Thus, we don't have strong intuitive evidence here *against* the claim that definite descriptions sometimes have a referential sense.

Given that there is a 'unitary' account, according to which uses of definite descriptions always semantically express what they mean on an attributive reading of them, and that accounts for their usage as well as does the thesis on which there is a semantic ambiguity here, Kripke argues that we should favor the unitary account on general methodological grounds (Kripke 1977: 18–20). Thus the lack of direct intuitive evidence on either side would leave the ambiguity thesis a loser, for Kripke—though he reports that making the determination just on such methodological grounds would leave him 'uneasy' (1977: 20). (Kripke's unease seems to be eased somewhat by the existence of a piece of direct evidence against the ambiguity thesis that he turns up at (1977: 21).)

This brings up a methodological question concerning the evaluation of contextualism that, despite its importance, I will mention only here in this note—largely because I don't have much of an answer to it: Is there a methodological presumption against typical contextualist theories that is like the one Kripke presses against ambiguity theories? (It's interesting to note that Kripke uses 'two senses of "know(s)"' theories as an example of the kind of ambiguity theories that should be avoided at (Kripke 1977: 19).) Or, in the way relevant to such presumptions, should typical contextualist theories be construed as unitary theories, according to which the terms in question have one sense, albeit a context-sensitive one, and thus do not run afoul of the general (though defeasible) advice against multiplying senses beyond necessity (cited at Kripke 1977: 20)? I would be happy enough to accept that there is a general presumption here in favor of invariantist over contextualist accounts. In that case, my arguments in this volume can be read as attempts to show that contextualism overcomes that presumption. But I'm far from sure that's the way to go, and actually suspect that typical contextualist

But this lack of clear intuitions seems connected with another point I've made in this chapter. The intuitions I endorse in the case for contextualism are cases in which speakers *who have no false beliefs about relevant underlying matters of fact* seem to appropriately, as well as to truthfully, make the assertion in question. The intuitions that such assertions are true, I've argued, deserve a special degree of trust because they are reinforced by the intuition that they are appropriate, which provides further reason to think they are true *in cases where the speaker has no underlying false belief*, because of the presumption that speakers cannot appropriately assert what from their own point of view (what, given their beliefs about the relevant underlying matters of fact) is not true (see sections 2 and 3, above). In cases where the speaker *does* have such an underlying false belief, the presumption only supports the conclusion that the speaker's claim is true *from her own point of view*, and doesn't get us as far as the ambiguity theorist needs here—to the destination that the speaker's claim is true. (Those who claim the speaker's assertions here, like all uses of definite descriptions, have the content assigned to them by a purely attributive theory can happily agree that the speaker's claim in this case is true *from his own point of view*. What they must resist is that the assertion is just true.) And in all the relevant cases of a supposed referential use of a definite description, it seems there is a crucial false belief the speaker has about an underlying matter of fact: In the case we have been considering, that mistake is that the man the woman is talking with is her husband. What we need is a speaker who will, with seeming truth and appropriateness, assert that 'Her husband is kind to her', even though the speaker has no such underlying false belief. But as Kripke emphasizes, that's precisely what we don't get here: Those of us who realize that the person who is actually her husband is not kind to the woman are not inclined to assert 'Her husband is kind to her', and if we were to assert that, our assertion would intuitively seem far from unproblematic.

accounts at least don't have as strong a presumption against them as do ambiguity theories. (In other work, I've found myself claiming that a contextualist theory of 'might' counterfactuals that I was advocating is to be preferred on methodological grounds over an ambiguity thesis championed by David Lewis (DeRose 1999b: 402)—though I also argued that there is direct evidence favoring the contextualist account. However, that claim was made against the special background of our having good reasons to posit the semantic flexibility utilized by the contextualist account in order to account for other phenomena, anyway.)

3

Assertion, Knowledge,
and Context

This chapter brings together two positions that for the most part have been developed and defended independently of one another: contextualism about knowledge attributions and the knowledge account of assertion.

The positions under discussion are both located in the area of overlap between epistemology and the philosophy of language, and they have both received a good deal of attention in recent years. But there is further reason for surprise that more has not been done to bring them together. The chief bugaboo of contextualism has long been the concern that the contextualist is mistaking variability in the conditions of warranted assertability of knowledge attributions for variability in their truth-conditions. This strongly suggests that a serious assessment of contextualism will demand a discerning look at the question of what it takes for a speaker to make a warranted assertion. And the knowledge account of assertion—according to which what one is in a position to assert is what one knows—promises to provide a large and important part of the answer to this question.

It also turns out that the knowledge account of assertion dissolves the most pressing problem confronting the positive cases for contextualism from ordinary language that we pursued in the previous chapter, and, on top of that, provides a second powerful positive argument in favor of contextualism. Or so I will argue.

It will take a bit of work to uncover the most pressing problem for the contextualist's argument from ordinary language, since critics of contextualism have not themselves been very proficient in identifying it. In sections 1–3, then, we will examine the 'warranted assertability objection'—the charge that the contextualist mistakes a variability in the conditions for the warranted assertability of knowledge attributions for a variability in their truth-conditions, concluding that, as it stands, that objection is powerless. But in section 4, I will propose a new version of the warranted assertability objection—the 'generality

objection'—that may be what is behind the worries of those who propose the warranted assertability objection, and that may well constitute contextualism's most pressing problem, and will explain why it is such a serious worry. In sections 5–9, we turn our attention to the knowledge account of assertion, the grounds that support it, its need for contextualism, and the positive argument for contextualism that it provides. In sections 10–11, we use the knowledge account of assertion to squash the Generality Objection and we also cast other shadows over the prospects for anti-contextualism.

1. The Classical Invariantist's Warranted Assertability Objection

In Chapter 2, we saw how contextualism can be argued for from an important fact about ordinary usage that is exemplified in the pairs of cases that contextualists appeal to: that in some contexts we are quick to call 'knowledge' what we will deny is 'knowledge' in other contexts.

Wisely, defenders of invariantism typically accept the evident facts about usual and, in the relevant sense, appropriate linguistic behavior that the contextualist's case pairs display: Indeed, that is how we talk, and appropriately so. What is controversial is whether these varying standards for when ordinary speakers *will* attribute knowledge and when they're *warranted* in attributing knowledge reflect varying standards for when it is or would be *true* for them to attribute knowledge, and the contextualist's claim is the more controversial one concerning a variation in truth-conditions.

Accordingly, the most influential source of resistance to contextualism is the *warranted assertability objection*—the charge that what the contextualist takes to be a variation in the truth-conditions of knowledge attributions is in reality only a variation in the conditions for the warranted assertability of those claims.[1] The warranted assertability objection, then, is a way of

---

[1] In one sense, I take Unger (1984) to be the 'birth' of contextualism: Though contextualist theories of knowledge were put forward before then (especially contextualist versions of the relevant alternatives theory), it was to my knowledge in that work that the terms 'contextualism' and 'invariantism' were first used as I am using them here, to label just these positions, and, more importantly, it was here that these rival positions were separated out from the accounts of knowledge in which they were usually embedded for special treatment, to receive their first careful expositions and assessments in their own right. (See my brief history of contextualism in section 12 of Chapter 1.) If we do mark Unger (1984) as the birth of contextualism, then the theory was born facing the warranted assertability objection. In that work, Unger was not endorsing contextualism, but defended the position that it was simply indeterminate whether it or invariantism was correct. And the main ground for Unger's verdict was that the invariantist's account of our ordinary use of 'know(s)', on which the varying epistemic standards

resisting one or the other of the key premises of the contextualist's main argument.

If the contextualist has chosen his pair of cases well, then, at least when the cases are considered individually, there will be a quite strong intuition about each of the assertions (the positive ascription of knowledge in LOW and the denial of knowledge in HIGH) that it is true. So the invariantist must deny a quite strong intuition—and an intuition that, for various reasons we looked at in Chapter 2, is quite solid. The warranted assertability objection is an attempt to explain away the considerable power of one or the other of two intuitions, which together are lethal to the invariantist. Perhaps the positive ascription of knowledge in LOW is really false, but seems true because the low standards *for warranted assertability* in place there make the positive ascription of knowledge quite appropriate, and we mistake the warranted assertability of the claim for truth. Or maybe it's the denial of knowledge in HIGH that is false but appropriate: Given the high standards for the warranted assertability of knowledge in place there, a positive claim that the subject knows would be unwarranted (though true), and it's the denial of knowledge that is appropriate (though false). The invariantist wielding the warranted assertability objection may often sensibly decline to say which of our intuitions is wrong. About some cases, she might admit that it's hard to say whether the subject knows or not. She will merely claim that the fact that a fairly clear appearance of truth attaches to both the positive ascription of knowledge and the denial of knowledge is due to the warranted assertability that both of the claims enjoy, each in its own context.

Note that the subject-sensitive invariantist will not have to deny either of the contextualist's premises regarding *first-person* knowledge claims, and will typically agree with the contextualist that both the assertion of 'knowledge' in LOW and the denial of 'knowledge' in HIGH are true in such first-person cases. But, as we saw in Chapter 2, the contextualist's best cases are third-person cases, and here the defender of SSI will join other invariantists in denying the contextualist's premises as applied to those third-person cases—and in needing a way to make that denial plausible. Here they *may* join the classical invariantists in wielding the warranted assertability objection against the contextualist's premises. However, that would seem to put subject-sensitive invariantists in a somewhat awkward position: If they utilize the warranted assertability

that govern its use are merely standards for the warranted assertability of the relevant claims, is just as good as the contextualist's account. While versions of the objection have appeared here and there in print since then, I have experienced it mainly as something that has been 'in the air' in philosophical circles (or at least in the philosophical circles I have traveled in): Raising this objection seems to be the immediate reaction to contextualism of many philosophers.

objection against the contextualist's premises concerning third-person cases, why are they not similarly skeptical about the truth-values of the first-person claims and denials of knowledge, where they agree with the contextualists? At any rate, the warranted assertability objection is typically used in defense of *classical* invariantism, in which case it is wielded uniformly against the contextualist's premises concerning both first- and third-person cases. The objection is an expression of the quite consistent attitude toward ordinary usage typical of classical invariantism: that the varying epistemic standards that apply to uses of 'know(s)' in different contexts govern only whether it is *appropriate or warranted* to *say* that someone does or doesn't 'know', while the epistemic standard for whether a subject really does know—whether it would be *true* to say that they 'know'—does not vary from context to context. Consequently, in this chapter our focus will be on contextualism's battle with classical invariantism; discerning between contextualism and subject-sensitive invariantism will be the aim of Chapter 7.

The warranted assertability objection to contextualism is an example of what we can call a *warranted assertability maneuver* (WAM). A WAM involves explaining why an assertion can seem false (or at least not true) in certain circumstances in which it is in fact true by appeal to the fact that the utterance would be improper or unwarranted in the circumstances in question. Going the other way, an intuition that an assertion is true can be explained away by means of the claim that the assertion, while false, is warranted, and we mistake this warranted assertability for truth. Either way, the maneuver is based on the correct insight that truth and warranted assertability (and also falsehood and unwarranted assertability) are quite different things, but that we can easily mistake one for the other.

Though that general insight is sound, and though some instances of this type of maneuver have been correct, there is a distinct danger that WAMs can be carried too far and misused, as can be illustrated by the below fiction.

2. The Myth of Jank Fraction: A Cautionary Tale

Suppose that, for some reason we needn't concern ourselves with, some philosophers toward the middle of the twentieth century took a liking to the notion that all it takes to know a proposition is that one believe it.[2]

---

[2] I don't know of any philosophers who have endorsed such a view. However, as we saw in Chapter 1, section 6, the idea that knowledge requires only *true* belief has been accepted.

We may suppose that their view came to be called the 'Equivalence Thesis' because, according to it, 'S knows that p' is equivalent to—has the same truth-conditions as—'S believes that p'.

You don't have to be a veteran of the Gettier wars to be immediately struck by counter-examples to this sorry theory. Some accounts of knowledge face the problem of Lucky Louie: They counter-intuitively count as knowers subjects who believe a proposition in some baseless way—say, by means of a wild guess—but, by sheer luck, turn out to be right. The Equivalence theorists faced not only that, but also the problem of Unlucky Ursula, whose belief is not only based on a pure guess, but is also false! Indeed, we may suppose that among the many counter-examples hurled at the hapless early Equivalence theorists, ones involving subjects with false beliefs were thought to be the most pressing.

Enter Jank Fraction, who defended the Equivalence Thesis against its most pressing problem by adding to the Thesis the auxiliary hypothesis that an assertion of 'S knows that p' generates an implicature to the effect that p is true, and thus a warranted assertability condition of 'S knows that p' is that p be true. This, however, is *only* an implicature and a warranted assertability condition of 'S knows that p', according to Fraction, not a truth-condition. Thus, according to Fraction, the Equivalence Thesis is correct about the truth-conditions of 'S knows that p'. Fraction argued that since his 'Supplemented Equivalence Thesis'—the Equivalence Thesis regarding the truth-conditions of knowledge attributions together with his additional claim about those sentences' warranted assertability conditions—predicts that 'S knows that p' will be unassertable where p is false, it can explain why we think such a knowledge attribution is false where the relevant p is false, and thus our intuitions that false beliefs don't amount to knowledge really don't hurt his theory. We may finally suppose that Fraction closed his 1979 defense of the Equivalence Thesis, 'On Assertion and Knowledge Attributions', with these words:

In my view this puts a very different complexion on certain putative counterexamples to the Equivalence thesis. We saw, for instance, how the Equivalence theorist must hold that 'Unlucky Ursula knows that Carter was re-elected' is true. But what is it that is *immediately evident* about this putative counterexample? Surely that it has very low assertability. But the Supplemented Equivalence theory *explains* this, and what a theory well explains cannot be an objection to that theory.[3]

---

[3] The character Jank Fraction is based on the actual Frank Jackson, and this passage is a simple transformation of the closing paragraph of Jackson (1979).

3. Lame WAMs and the Warranted Assertability Objection to Contextualism

There's something fishy about Fraction's maneuver, as I hope you can sense. Fraction's defense should not be allowed to mitigate the negative verdict we are inclined to reach against the Equivalence Thesis. The sense that this WAM is unsuccessful is partly due to the fact that it is being offered in defense of such a loser of a theory—a theory for which it is difficult to imagine what positive support it might have.[4] But it is important to notice the deeper reasons why this defensive maneuver should be allotted no force here. For now, we'll focus on this important reason: Fraction's WAM is an instance of a general scheme that, if allowed, could be used to explain away far too easily the counter-examples marshaled against *any* account of the truth-conditions of sentences in natural language. Whenever your theory seems to be wrong because it is omitting a certain truth-condition—as Fraction's theory of 'S knows that p' seems to wrongly neglect to include a condition to the effect that p must be true for a sentence ascribing knowledge of p to a subject to be true—you can simply claim that assertions of the sentences in question generate implicatures to the effect that the condition in question holds. Thus, you claim, the alleged condition is a warranted assertability condition of the relevant sentences: to assert such a sentence where the condition does not hold will be unwarranted because the speaker will generate a false implicature. Your critic, you claim, is mistaking the falsity of an implicature that is generated by assertions of the sentences in the relevant circumstances for the falsity of what the speaker says in uttering the sentence, and is thus mistaking what is only a warranted assertability condition of the sentence for a truth-condition of it.

So armed, you are ready to defend, say, an account of 'S is a bachelor' according to which S's being unmarried is not a truth-condition of the sentence. Of course, real and imagined married men will provide a slew of apparent counter-examples to your besieged theory. But rather than *correct* your account to remedy your apparent mistake, you can follow Fraction's example

---

[4] By contrast, the actual philosopher Frank Jackson, on whom my fiction of Jank Fraction is based, utilized a maneuver in some ways similar to Fraction's, but in defense of a theory for which there are genuine positive reasons. The 'Equivalence' theory that the actual Jackson defends claims that indicative conditionals are equivalent to material conditionals (that is, they are true iff their antecedent is false and/or their consequent is true). One reason (the best reason, to my thinking, though not to Jackson's, he has informed me) for accepting *that* Equivalence theory is to be found in the apparent validity of arguments of the form '*P or Q; therefore, If not-P, then Q*'. Example (due to Robert Stalnaker): 'Either the butler did it or the gardener did it. Therefore, if the butler didn't do it, then the gardener did' (Stalnaker 1975: 269).

by stubbornly holding fast to your proposal about the truth-conditions of 'bachelor' sentences, and *supplementing* your theory with the auxiliary claim that S's being unmarried is an implicature generated by an assertion of 'S is a bachelor' and is therefore a warranted assertability condition for the sentence, though not a truth-condition. Thus, your supplemented theory explains why 'S is a bachelor' seems false when said of a married man. And what a theory well explains …

'But wait! Aren't there WAMs that we rightly do give credence to?'

Yes. Here's one that I myself have used (DeRose 1999a: 196–7)—I believe effectively. When a speaker knows that p, it can seem somehow wrong, and to some it may seem downright false, for him to say 'It's possible that p_{ind}.'[5] Suppose, for instance, that Ringo wants to borrow a certain book, and he asks Paul whether the book is in Paul's apartment. If Paul knows full well that the book is there, it would be somehow wrong for him to answer, 'It's possible that it's there.' Indeed, pre-theoretically, many feel some tendency to judge that Paul would then be saying something false. (Many will find this tendency in competition with an opposing intuitive pull toward the verdict that the statement is true—though somehow misleading and wrong.) Such tendencies could tempt one toward a 'Don't Know Either Way' (DKEW) account of 'It's possible that p_{ind}', according to which:

> DKEW: S's assertion, 'It's possible that p_{ind}' is true iff (1) S doesn't know that p is false and (2) S doesn't know that p is true.

But this temptation should be resisted, I believe, for the correct account lies down the simpler 'Don't Know Otherwise' (DKO) path:

> DKO: S's assertion, 'It's possible that p_{ind}' is true iff S doesn't know that p is false.[6]

According to DKO, Paul is asserting the truth in our example. To the extent that one does have the intuition that Paul is saying something false, the backer of DKO can attempt to explain away that intuition, and other intuitions to the effect that speakers who know that p speak falsely in saying that p is possible, as follows. Both 'p' and 'I know that p' are stronger than—they imply but are not implied by—'It's possible that p_{ind},' according to DKO.[7] And there's a

---

[5] The subscript 'ind' indicates that the embedded p is to be kept in the indicative mood: very different possibilities are expressed where the embedded p is subjunctive.

[6] Though this is the right path, DKO is not correct in detail as it stands. For my best attempt to get the details right, see DeRose (1991). Those details don't affect the current argument.

[7] As Patrick Rysiew points out (Rysiew 2005: 57), the 'Assert the Stronger' rule we are about to encounter is perhaps best construed as operating with a notion of relative 'strength' such that p can

very general conversational rule to the effect that when you're in a position to assert either of two things, then, other things being equal, if you assert either of them, you should assert the stronger of them (Grice 1961: 132; Jackson 1979: 566). When someone like Paul knows that p, they're in a position to assert that p—and they're often in a position to assert even that they know that p. Thus, by the 'Assert the Stronger' rule, they should assert one of those stronger things rather than the needlessly weak 'It's possible that p_{ind}.' To assert that weak possibility statement is then unwarranted and generates the false implicature that the speaker doesn't know that p by the following Gricean reasoning, which is based on the assumption that the speaker is following the 'Assert the Stronger' rule: 'If he knew that p, he would have been in a position to assert something stronger than "It's possible that p_{ind}," and thus would have asserted some stronger thing instead. But he did assert "It's possible that p_{ind}," not anything stronger. So he must not know that p.' We can easily mistake the falsehood of the implicature generated by such an assertion of 'It's possible that p_{ind}' and the resulting unwarrantedness of the assertion for the falsehood of the assertion itself.

Or so the backer of DKO can claim. And plausibly so. My purpose here is not to argue for the ultimate effectiveness of this WAM, or for the correctness of the DKO approach (though the maneuver is effective and the approach is correct), but to use this WAM as an illustration of one that, unlike the ones we considered above, has at least some force, as I hope the reader can already sense. We will later (in section 11) see further reasons for thinking that this WAM is quite credible. But for now, note this central reason for taking it seriously—at least more seriously than Fraction's lame maneuver and our imagined WAM used in defense of the strange theory of 'bachelor'. A vitally important contrast here involves *how* the wielders of these WAMs explain the

sometimes count as being 'stronger than' q, even though p doesn't entail that q. (This construal will likely be something along the intuitive lines of 'being more informative than'.) I am here citing that both 'p' and 'I know that p' imply but are not implied by 'It's possible that p_{ind},' as sufficient, but not necessarily necessary, conditions for their being stronger than 'It's possible that p_{ind}.' However the notion of relative strength is best worked out for the purposes of this rule, the comparisons of relative strength I'm depending on here should hold.

However, I think the operative rule still has to be construed along the lines I have in the text, even if the notion of strength at work is fairly subtle, and I don't think the rule can be replaced in the way that Rysiew (in discussion with Kent Bach) suggests: 'Better, then, to see the examples just given as illustrating Grice's maxim of quantity—Make your [conversational] contribution as informative as is required (for the current purposes of the exchange)' (Rysiew 2005: 57). The problem with the replacement is that it doesn't explain many of the cases that are well explained by 'Assert the Stronger'. For even when a weaker statement (e.g. 'It's possible that the book is in my office') is informative enough for the purposes of a given conversation (as can happen, for example, when it would have been fine for a speaker not to say anything about the matter in question), by asserting that weaker statement, one *still* often implicates that one isn't in a position to assert a stronger statement one could have easily made if one had been in a position to assert it.

generation of the implicatures to which they appeal. The defense of DKO utilizes a *general* rule of conversation—'Assert the Stronger'—which applies to assertions of any content.[8] This general rule, together with DKO's account of the content of 'It's possible that p_{ind}', generates the implicature that S doesn't know that p, by reasoning of a familiar, Gricean style. By contrast, Fraction's WAM and the defense of the crazed theory of 'bachelor' each resort to positing a *special* rule attaching only to assertions involving the relevant term—'knows' or 'bachelor'. This is what rightly arouses our suspicions that any theory that omits what is in fact a truth-condition for a type of assertion could just as well execute this maneuver. But it's not so easy to generate the implicatures you need to deflect the apparent counter-examples to your theory by means of *general* conversational rules that can be tested on very different sentences. If your theory of the truth-conditions of the relevant sentences, together with such general rules that are independently motivated, really does predict the occurrence of apparent counter-examples to your theory, that would seem legitimately to mitigate the force of those apparent counter-examples.[9]

---

[8] Don't mistake generality for exceptionlessness. Though the rule is very general, there are occasions on which one should not assert the stronger, as is pointed out by Jackson (1979: 569–73).

[9] This preference for explanations generated by general rules over those that utilize special rules amounts to, in Grice's (1989) terminology, a preference for uses of conversational implicatures over conventional implicatures. I join Kent Bach (1999) and Jason Stanley (2002) in being wary of conventional implicatures, but the character of my wariness is more like the wariness that Grice himself displayed. I am largely sympathetic to Stanley's statement (made appealing to the arguments in Bach) that 'the category of conventional implicature is simply suspect. The arguments for the existence of conventional implicature rest entirely on rather unconvincing appeals to intuitions. Conventional implicature is also theoretically artificial, and sits uneasily with other Gricean distinctions' (2002: 333). However, some intuitions that support the existence of conventional implicatures strike me as worthy of some credence, even if they are not compelling. I have in mind here intuitions like those that support the favorite example of conventional implicature: that 'p but q', while equivalent to 'p and q' in its truth-conditions, does generate a conventional implicature, not carried by 'and', that there is some contrast between p and q (see, for example, Jackson 1979: 573–4). I am at least open to this being a genuine example of conventional implicature (but see Bach's discussion), so my skepticism is not really over the existence of conventional implicatures. Rather, my worry is that even if some conventional implicatures do exist, it is problematic to posit a conventional implicature to 'grice away' intuitions that run counter to a semantic thesis one is defending, as we're imagining Fraction is attempting to do. This is like the wariness of Grice, who accepted some cases as being examples of conventional implicatures, but who, after displaying a supposedly model case of conventional implicature—the case of 'but'—rather warily writes, 'In any case, the nature of conventional implicature needs to be examined before any free use of it, for explanatory purposes, can be indulged in' (1989: 46).

Note that even if 'but' provides us with an example of conventional implicature, our access to this conventional implicature is through a 'strange but true' intuition: We find that, where p and q are both true but there is no relevant contrast between them, 'p but q' is somehow wrong to assert, but still true. (Example: ? 'The Cubs scored no runs at all, but they lost the game.') These 'strange but true' intuitions, which are needed to establish the existence of conventional implicatures in the first place, would at the same time undermine the methodological presupposition behind Fraction's—and Jackson's—procedure: namely, that our immediate access is only to warranted/unwarranted assertability,

The problem anti-contextualists face—or, as we will see later, just *one* of the daunting problems they face—as we turn back to the warranted assertability objection against contextualism, is that, as this objection has so far been formulated, it does no such thing. Defenders of invariantism have not proposed credible accounts on which general rules of conversation together with their proposed invariantist accounts of the truth-conditions of knowledge attributions predict the pattern of varying conditions of warranted assertability of knowledge attributions that we encounter in natural language. (A pair of invariantist accounts that can appear to be exceptions to this are discussed in the appendix to this chapter.) Truth be told, the warranted assertability objection against contextualism often takes the form of a *bare* warranted assertability objection: It's simply declared that it's the conditions of warranted assertability, rather than of truth, that are varying with context, and the contextualist is then accused of mistaking warranted assertability for truth. To the extent that defenders of invariantism go beyond such bare maneuvers, their hints tend to point in the direction of *special* rules for the assertability of 'knows', such as 'If someone is close enough, for present intents and purposes, to being a knower, don't say that she doesn't know, but rather say that she knows.' And of course, if he's allowed to appeal to the bare possibility that warranted assertability is being confused with truth, or to special rules about the term in question, even our theorist about 'bachelor' can rebut the evidence against his theory.

Indeed, as we'll see in section 11, if we consider contrasts between forceful and forceless WAMs different from the one we have just contemplated and attempt to discern various features WAMs should possess in order to be accorded any mitigating force, and we then apply these results to the case at hand, we find that the warranted assertability objection against contextualism fails miserably to meet all manner of other reasonable criteria we can discern for what it would take for a WAM to be successful.

4. The Generality Objection

The contextualist may simply take this all as good news and declare victory, as I did in earlier work.[10] But one can't help but think that we haven't gotten to

and our intuitive access to truth/falsehood is only through warranted/asserted assertability (see the imagined Fraction quotation at the end of section 2, above, and the actual Jackson passage, on which Fraction's is modeled: 1979: 589). For if our only intuitive access to truth/falsehood came indirectly through warranted/unwarranted assertability, how could we get the 'true' part of 'strange but true' intuitions here, where the sentences are unwarranted?

[10] DeRose (1999a); see esp. section 11.

the bottom of things yet. For very smart philosophers, who would not be at all tempted by our imagined lame defense of the crazed theory of 'bachelor', nevertheless do find the warranted assertability objection to contextualism quite attractive, and I must confess that I can feel the attraction, too, though I ultimately want to reject the move. Can it really be *such* a loser of a maneuver? What is it about that *particular* appeal to warranted assertability considerations that makes it so attractive?

An explanation for that appeal is not that far to seek, especially when we keep in mind the concerns about generality just raised. Those concerns should lead us to ask: Is the phenomenon we're dealing with particular to knowledge attributions, or do the varying epistemic standards we've discovered affect other, quite different, assertions as well?

Here the invariantist should be pleased to note that this variation is ubiquitous, affecting the standards, at least for warranted assertability, of assertions whose truth-conditions obviously don't vary as epistemic standards change. Indeed, varying standards do in some way govern our use of 'S knows that p', but what about that embedded 'p', which can be just about any proposition? Are not the 'high-standards contexts' in which it becomes very difficult to warrantedly assert 'S knows that p' also contexts in which it becomes wrong to assert the simple 'p'? In the high-standards Bank Case, for instance, where it seems I can't claim to know the bank will be open on Saturday, it also seems it would be wrong for me to assert, flat-out, 'The bank will be open on Saturday.'

In the case of first-person knowledge attributions—'I know that p'—the assertability of the knowledge claim and the assertability of the simple 'p' seem to fade away together as the epistemic standards governing a conversation become more demanding. The knowledge claim may lead the way in the move toward unassertability: It will seem more problematic than the assertion of 'p' in various intermediate contexts. But the simple 'p' follows closely enough behind that it's difficult to come up with cases in which 'I know that p' has become clearly unwarranted while 'p' is still clearly warranted. Things are quite different with second- and third-person knowledge attributions. Here, of course, the speaker will often be in a position to assert 'p' but will clearly be in no position to assert that a certain S knows that p—sometimes simply because the speaker doesn't know anything about S. Still, if we start with a second- or third-person case in which 'p' and 'S know[s] that p' actually are both warranted, and then raise the standards in our imagination, we find that the two claims move toward unassertability together—though, again, the knowledge attribution may lead the way.

These observations give rise to a serious challenge to contextualism that I will call the *Generality Objection*.[11] There is some very general rule of conversation to the effect that one should assert something only if one is positioned well enough with respect to that proposition to properly assert it, and thus a condition on warranted assertability that one be epistemically positioned well enough with respect to what one asserts.[12] We may of course hope for a more specific rule along these lines—one that specifies how well positioned one must be with respect to a proposition to be able to assert it. (And we are about to discuss such a more specific rule in the following section.) But at this point, the Generality Objector is seeking to defend an invariantist account of the context-variability in the assertability of knowledge attributions by appealing only to the fact that there is *some* such requirement on warranted assertion and to the observation that in general, and not just with respect to knowledge-ascribing sentences, how well positioned one must be with respect to something to be able to properly assert it is a context-variable matter. For as the Generality Objector will correctly point out, when the simple 'p' becomes unassertable because we are moving into more demanding contexts, this is generally not due to any change in the truth-conditions of 'p'. 'p' can be just about anything, and most of our assertions are obviously insensitive in their content to what epistemic standards happen to be governing their use. For example, though the content of an assertion of 'The bank is open on Saturdays' may be context-sensitive in various *other* ways, the conditions under which it is *true* (as opposed to assertable) clearly don't depend at all on what epistemic

---

[11] Great thanks to Paul Boghossian for first raising (in conversation, circa 1992) a version of this very insightful objection. The details of how I develop it here, of course, are not his responsibility.

[12] One might question whether assertion is governed by a rule concerning how well positioned one must be with respect to a proposition in order to be able to properly assert it. What any such rule seems ultimately aimed at securing is that, at least often enough, only the truth is asserted. Thus, in the simplest scenario—and perhaps therefore the scenario we should presume, absent any reason to think matters are more complicated—assertion would not be governed by any rule about how well positioned one must be in order to assert, but rather by the simple rule most directly tied to that ultimate aim: Assert only what's true. Thus, Timothy Williamson, in defending his version of the knowledge account of assertion, expends much effort in battling the rival 'truth rule', according to which one must assert only what is true (2000: 244–9). See also n. 16, below.

The rule here—that one should assert only that with respect to which one is positioned well enough to assert—is basically Grice's second maxim of Quality—'Do not say that for which you lack adequate evidence' (1989: 27)—correcting for Grice's overly narrow focus on evidence. I think a speaker should be counted, in Gricean terminology, as having generated a conversational implicature to the effect that she is in a position to assert that p when she does assert that p, though Grice himself might have declined to call this an 'implicature' at all, since it simply specifies that the speaker is obeying the second maxim of Quality, and Grice stipulates that 'implicature' excludes commitments that are *trivially* generated by the supposition that the speaker is observing the Cooperative Principle and its associated conversational maxims (1989: 41–2). For discussion of whether such commitments should be counted as implicatures, see DeRose and Grandy (1999: 417–19 nn. 13 and 19).

standards are in play when the sentence is asserted. But, as is generally agreed, the truth of 'p' is a necessary condition of the truth of 'S knows that p'. Thus, it is the reverse of surprising that 'S knows that p' should become unassertable due to high epistemic standards when (or roughly when) 'p' thus becomes unassertable: If one is not positioned well enough to assert that 'p', then of course one won't be positioned well enough to assert the stronger 'S knows that p'. Since 'p' becomes unassertable in high-standards contexts even though there is no change in its content as we move into high-standards contexts, and since the drift toward the unassertability of 'S knows that p' as we move into more demanding contexts is just what we would expect given that 'p' displays a similar drift, why suppose that the unassertability of the knowledge claim in high contexts is due to a change in content *it* undergoes as we move into such contexts? According to the Generality Objection, there is no good reason to suppose there is such a variation in truth-conditions of knowledge attributions.

That's a tough question, and a tough objection. To answer it, as I will attempt to do in section 10, I will employ the knowledge account of assertion, to which we now turn.

5.   The Knowledge Account of Assertion

What is the content of the rule that prohibits us from asserting what we are not positioned well enough to assert? How strong a position must one be in with respect to a proposition to be able properly to assert it? A little reflection on these questions can quickly lead to pessimism about there being any sufficiently general answer that is any more informative than 'Assert only what you are positioned well enough with respect to', since it's fairly clear that, as the Generality Objector observes, one is here shooting at a moving target: In some conversational contexts, one must be extremely well positioned with respect to a proposition to be able to properly assert it, while in other contexts one can properly assert something with respect to which one is only moderately well positioned.

Nevertheless, an impressive answer to our question has emerged, and there are good reasons to think it is correct. Appropriately enough, according to the *knowledge account of assertion*, one must *know* that p in order to be positioned well enough with respect to p to assert it.

The knowledge account of assertion has been packaged in two basic forms. Following the basic direction of comments by G. E. Moore, most advocates

have tended to formulate the account in terms of a principle to the effect that when one asserts that p, one represents it as being the case that one knows that p.[13] But more recently, Timothy Williamson, in defending the knowledge account, has expressed it as a claim, not about what one represents as being the case in making an assertion, but about what rule governs the practice of assertion: According to Williamson, the constitutive rule of the practice of assertion is that one should assert only what one knows.[14] For our current limited purposes, these are just two sides of the same coin: If one represents oneself as knowing that p by asserting p, then, to avoid falsely representing oneself, one should follow the rule of asserting only what one knows; and if assertion is governed by a rule that one should assert only what one knows, then one will represent oneself as knowing that p when one asserts that p. We can leave it open whether to follow Williamson in holding that this rule is *the* single constitutive rule specific to the practice of assertion, *all* other rules governing assertion being consequences of this single rule, together with more general norms not specific to assertion.[15] I *will* join Williamson in holding that this is the only rule governing assertion that has to do with asserting only what one is positioned well enough with respect to—or, when using the other form of the knowledge account of assertion, that the strength of the position that one represents oneself as being in when one asserts that p is just that of knowing that p, nothing more nor less. What we leave open here is whether all of the other, quite different rules of assertion that don't have to do with asserting only what one is positioned well enough with respect to—like those enjoining us to assert only what's conversationally relevant—can be derived from the knowledge rule together with other rules not specific to assertion. In either of its forms, then, the knowledge account of assertion says that one is positioned well enough to assert that p iff one knows that p.

As happens with other rules, a kind of secondary propriety/impropriety will arise with respect to this one. While those who assert appropriately (with respect to this rule) in a primary sense will be those who actually obey it, a speaker who broke this rule in a blameless fashion, by asserting something

---

[13] See Unger (1975: chapter 6), Slote (1979), and DeRose (1991: 597–605). Williamson (2000: 252) reports that Lloyd Humberstone suggests that this talk of 'representing oneself' in certain ways by making assertions may have begun with Max Black, who wrote: 'In order to use the English language correctly, one has to learn that to pronounce the sentence "Oysters are edible" in a certain tone of voice is to represent oneself as knowing, or believing, or at least as not disbelieving what is being said' (Black 1952: 31).

[14] See Williamson (2000: chapter 11), which updates Williamson (1996).

[15] (Williamson 2000: 241). That there is such a unique rule is little more than an item of faith for Williamson, with no justification offered other than that a simple account consisting of such a single rule would be 'theoretically satisfying' (2000: 242).

she didn't know but reasonably thought she did know, would in some secondary sense be asserting properly, and a speaker who asserted something she reasonably thought she did not know, but in fact did know (if this is possible), would be asserting improperly in that secondary sense. An example of someone reasonably but falsely thinking they know something is provided by Ginet's famous 'barn case' (see n. 19 of Chapter 1), which can be (and in fact most naturally is) imagined so that the character is reasonable in thinking that he knows that what he's looking at is a barn: Henry quite reasonably thinks that he knows that he's seeing a barn, because he reasonably thinks he is in a very normal situation; unbeknownst to him, he's in a region teeming with fake barns and in fact doesn't know, though in fact he happens to be seeing a real barn right now. In many other cases, a speaker reasonably thinks she knows, though in fact she doesn't, for what she believes is not even true. The knowledge account of assertion would lead us to expect that though such speakers are breaking the rule for assertion, they are blameless for doing so and their assertions are warranted in a secondary way, since they reasonably take themselves to know what they assert. Thus, our sense that such speakers are at least in some way asserting appropriately and are not to be blamed for their assertions does not falsify the knowledge account of assertion, which would lead us to expect just that. I believe the knowledge account fits the relevant data, and, in fact, provides the best account of the full range of data, though I am not here pursuing a full-fledged defense of the account. What would show that the knowledge account of assertion is too strong are cases in which it's apparent that a speaker properly asserts what she doesn't even reasonably take herself to know.[16] It would also be trouble for the knowledge account if we had a clear sense that there is nothing at all amiss in a speaker's asserting something she does not know in the cases where she reasonably takes herself to know what she asserts. But I certainly lack any such clear sense; in fact, I instead have a *fairly* clear sense that something *is* amiss in such cases—though, of course, the exact nature of what has gone wrong is open to a lot of negotiation. And, as in other cases where an agent violates a norm in a blameless fashion because

---

[16] Similarly, the truth account of assertion (according to which assertion is governed by the 'truth rule', mentioned above in the first paragraph of n. 12) is not refuted by the existence of just any example in which a speaker asserts a falsehood, but we can sense that she is in some sense warranted in making her assertion, for it may be that the speaker reasonably believes that she is speaking the truth. What would, and do, spell trouble for the truth account (by showing it to be too weak) are cases in which a speaker is not warranted in asserting some p that both is true and is such that she very reasonably takes it to be true. In many such cases, the speaker is in a position to say something like 'P is very likely,' 'It's almost certain that p,' 'I believe that p,' or even, 'I reasonably believe that p,' but, as we can all sense, is in no position to flat-out assert the simple 'P'. (See also Williamson's discussion of the truth rule, cited above in n. 12.)

she reasonably takes herself to be acting so as to be in line with the norm, there seems to be a certain kind of responsibility that the speaker bears for what has gone wrong—a kind of responsibility that is consistent with the way in which the speaker nonetheless acted blamelessly in asserting what she did.[17] Even in these cases, to acknowledge that one didn't know what one asserted seems to be a kind of admission and can be a type of taking responsibility for something that seems to have gone wrong in one's act of asserting—though in a way that's compatible with, and can be conjoined with, a claim that (or explanation for why) one was in some way acting reasonably and blamelessly in so asserting.

The case for the knowledge account of assertion has been (powerfully) made elsewhere,[18] and I will not here rehearse that whole case, or even any

[17] Igor Douven's argument against me here (and also against a related move by Williamson) at Douven (2006: 476–80) *seems* to be based on a sense that Douven seems to have that whatever goes wrong in such cases (Douven focuses on cases in which the speaker asserts something false, but she reasonably believes what she was asserting) is something for which the speaker in no sense bears any responsibility, though I'm far from certain that I'm understanding his argument correctly. At least with respect to the cases he explicitly treats, where the speakers have asserted something false, Douven not only recognizes that 'the assertion is wrong because its content is false', but also that 'this reflects negatively on the asserter, for it means that she is mistaken' (2006: 476). But this is combined with a claim that the speaker, at least in some way, is free (of responsibility?) for this: 'So I take the data to be explained to be, first, that there is something wrong with a false assertion but, second, that the asserter goes free if she had excellent reasons to believe that her assertion was true' (2006: 477). This data is then wielded against the knowledge account. However, as we've just seen, and as Douven seems to be aware, the knowledge account would lead us to expect that there would be a way in which the asserter would be blameless ('goes free' *of blame*) for her wrong assertion—at least in cases where she not only reasonably believed that what she asserted was true, but also reasonably believed that she knew it to be true. It isn't clear to me exactly why Douven thinks that the way in which such speakers 'go free' is not one that the knowledge account can handle—though my best guess here is that he's taking it as a datum that there is *no good sense* in which the speaker bears any responsibility for what has gone wrong. In that case, I disagree with him about that. It seems to me that one thing we do when we assert that p is take responsibility for p's being true, such that if it turns out that p is false, we are subject to complaints. In responding to such complaints, it seems that we can often legitimately cite the facts in virtue of which we should not be blamed for getting the matter wrong (facts in virtue of which we were reasonable to take ourselves to know and reasonable to assert that p), but this coexists with there being some good sense in which something for which we are responsible—something for which we have taken responsibility—has gone wrong. And, related to that, I also guess that Douven would say that there is nothing at all amiss—or at least nothing at all for which the speaker in any good sense bears any responsibility—in cases where the speaker reasonably took herself to know, but in which, though what the speaker said turned out to be true, it turns out that she didn't know what she asserted. In that case, I'm strongly inclined to disagree with him about that, as well, for, as I'm just about to relate, in such cases, it seems to me quite appropriate for a speaker to take responsibility for something that has gone wrong.

[18] While I still like and recommend the argument in Unger (1975: chapter 6), Williamson (2000: chapter 11) is now the state of the art. Williamson discusses and expands on some (but not all) of the grounds given by Unger and Slote (252–5), adds further positive arguments of his own, answers objections to the account (255–60), and argues against rival accounts. An important rival account is the probability account, on which one should assert only what is (in the relevant sense) sufficiently

part of it in any great detail. I will, however, quickly point out what to my thinking is one of the most important recommendations of the account: that it provides a nice handling of the knowledge version of Moore's paradox and other troubling conjunctions. Famously, Moore noted the oddity of assertions of the form 'P, but I don't believe that P'—this is the standard version of Moore's paradox. Less famously, but no less insightfully, Moore also noted the oddity of

1. P, but I don't know that p.

Moore's example: 'Dogs bark, but I don't know that they do' (1962: 277). Such conjunctions are clearly consistent: Of course p could be true without my knowing it! Yet they 'clash'; they sound like contradictions. Moore's handling of this clash is built on the claim that (as it was put by later writers in the Moorean tradition) in asserting that p, one represents oneself as knowing that p, or, as Moore himself put it: 'By asserting *p* positively you *imply*, though you don't assert, that you know that *p*' (1962: 277). Thus, when one asserts the first half of Moore's conjunction, one is representing it as being the case that one knows that dogs bark. Consequently, when one goes on to say in the second half of the sentence that one does *not* know that dogs bark, one is saying something inconsistent with what one represented as being the case in asserting the first half. This explanation is attractive because it supports our sense that *some* inconsistency is responsible for the clash involved in asserting the conjunction,[19] while, at the same time, happily removing that inconsistency from the realm of what's asserted: The conjunction asserted

probable. V. H. Dudman uses lottery cases to forcefully attack the probability account in his (1992), though he goes too far, I think, in his further claim that 'Assertability goes out of the window as soon as the underlying thought is reduced to relying on "mere" probability' (205); for discussion, see section DeRose (1996a: 576–8).

[19] I take this to be a very important advantage of this account over those that would merely account for why these conjunctions seem in some way odd, without accounting for why they 'clash'—why they sound like contradictions. After giving such a weaker account, at his (2006: 473–5), based on an account of why such conjunctions would be rarely used together with the claim that constructions will tend to sound odd to those who have not frequently encountered them, as often happens when people are learning a new language, Douven considers just the right pressing objection to his account (which he credits to Jonathan Adler): 'It might be objected to the foregoing that, by taking instances of [φ, but I do not know that φ] merely to sound odd, I am not doing full justice to the data, for these sentences sound odd in the very specific sense that they are heard as contradictory. And it seems at least far less straightforward to explain away a heard inconsistency in the pragmatic way that was just attempted than it is to explain away the mere odd-soundingness of a sentence. After all, it seems plausible that things we have never before encountered may, upon a first encounter, sound odd—but why should they sound contradictory? Surely the words and grammatical constructions of foreign languages we learned never sounded that way' (2006: 475–6). In reply, Douven cites an admission I made in DeRose (1991); here is Douven's complete reply to the objection, which immediately follows the material quoted above: 'However, I submit that what Keith DeRose (1991, 597) says when considering an instance of

is itself perfectly consistent, but in trying to assert it, one gets involved in a contradiction between one thing that one asserts and another thing that one represents as being the case.[20] And the knowledge account can also be used to explain other troubling conjunctions that I can't see how to handle

[φ, but I do not know that φ] goes for all of us: "I don't have a special feeling for inconsistencies; I can sense some kind of clash, but cannot distinguish my sensing of an inconsistency from my sensing of whatever it is that's wrong with the Moorean sentence." And as I hope to have shown in the preceding paragraphs, the sensed wrongness of Moorean sentences ("whatever it is") is open to an explanation that is compatible with the rational credibility account being correct' (2006: 476). But this misconstrues my admission. The admission is made immediately after presenting the Moorean sentence, reporting that it sounds like it is inconsistent, but agreeing that it is in fact consistent. The admission is that my ability to sense 'clashes'—apparent contradictions—does not reliably distinguish genuine from merely apparent contradictions. That doesn't mean that we don't have to account for why Moorean sentences 'clash', and that all we have to do is give any account which would explain why they seem odd in some way or other. In fact, explaining why these sentences 'clash'—seem inconsistent—is precisely the task I set for myself at (1991: 597), and take up at (1991: 597–8). And I am clear about the advantage of accounting for our sense of some kind of inconsistency: 'Part of the power of this explanation is that it supports our sense that some inconsistency is responsible for the clash. According to this explanation, an inconsistency *does* produce the clash, but it is *not an inconsistency in what one asserts* by conjoining the two halves of the sentence (a genuine inconsistency); it is rather an inconsistency between something one asserts and something one *represents as being the case* by asserting something else' (1991: 598).

 [20] At his (2005: 237–8), Matthew Weiner tries to account for why Moore's 'Dog's bark, but I don't know that they do' is paradoxical, without using the knowledge account of assertion, which he opposes, by claiming that one must have some warrant for the proposition that dogs bark to be in a position to assert the sentence, but that the most likely sources of warrant one might have would typically allow one to know that dogs bark. Weiner continues: 'If the speaker's situation is so unusual that she has some other warrant for believing that dogs bark, she should say so' (2005: 238). I don't think this adequately accounts for the clash of the specific Moorean sentence Weiner is dealing with, but, in any case, it *certainly* can't account for such clashes generally. 'P, but I don't know that p' quite generally produces the type of clash we are seeking to account for. But it cannot generally be presumed that whatever grounds a subject has for reasonably believing the p in question should also allow that subject to know that p. Weiner suggests that in cases where one's warrant for accepting a proposition couldn't be safely assumed to be such as to make one a knower—he thinks cases of predictions and retrodictions ('inferences to past unobserved happenings') typically fall in this category—one *can* properly conjoin an assertion of p with an admission that one doesn't know that p (2005: 238). But his example to show this—

I don't *know* they'll attack at nightfall—we haven't intercepted their orders—but my prediction is that they will

—seems to me to cheat in important ways. First, there is the matter of his emphasizing 'know': In checking for 'clashes' in conjunctions, I think it is crucial to *avoid* any strong or unusual emphasis, stress, or intonation, and stick with 'flat-footed' assertions of the conjunctions, for reasons explained at DeRose (1998: 70–2). But even more importantly, the second conjunct should just be 'they will [attack at nightfall]', rather than 'my prediction is that they will [attack at nightfall]'. While many predictions are assertions, that doesn't mean that 'my prediction is that p' must be an assertion of p; it as an assertion of what the speaker's prediction is, rather than being itself an assertion of p. The clean and fair presentation of Weiner's example—

I don't know that they'll attack at nightfall, but they will

—I submit, seems about as paradoxical as is the result of reversing the order of the conjuncts in Moore's own 'Dogs bark, but I don't know that they do': 'I don't know that dogs bark, but they do.' (Taking

without the account.[21] These and other grounds (see the works cited in n. 14, above) give us strong reason to think that the knowledge account of assertion is correct.

6. The Knowledge Account of Assertion Contextualized

How can we reconcile the correctness of the knowledge account of assertion with our earlier reason for being pessimistic about any such an account being

things the other way, if we are allowed to doctor up Moore's conjunction à la Weiner, it too ceases to produce a sharp clash: 'I don't *know* that dogs bark, but my claim is that they do.')

While I am discussing Weiner's paper, I will briefly take the opportunity here to register my reactions to his proposed counter-examples to the knowledge account of assertion at (2005: 230–1). These are proposed as cases in which the assertions are proper, though the speakers don't know the propositions they are asserting to be true. But it seems to me very far from clear that the speakers don't know the propositions they are asserting in these examples. Indeed, it wouldn't seem to me at all unusual for the speakers to have instead outright said, 'I know that...' in these examples. And where the speakers do refrain from claiming knowledge, making do with mere flat-out assertion in these examples, it would be perfectly appropriate for their conversational partner to ask, 'How do you know that?' If the speaker then admits that she doesn't know, it seems that she is then barred from making the flat-out assertion. (On contextual analysis, it seems likely that the epistemic standards have been raised in such cases.) At any rate, where one sticks to examples where it actually is fairly clear that the speaker doesn't know the proposition she's asserting, I think one will find it hard to come up with cases where the assertion is also clearly appropriate. At least, I haven't encountered any counter-examples that I've found convincing.

[21] In other work, I've used the claim that in asserting something one represents oneself as knowing it, together with independently supported analyses of the relevant modal statements, to also explain these other 'clashes' (all given here in general schematic form, except for 4, which is an example of the type of conjunction, involving a 'simple "might"', that I mean to indicate):

2. It's possible that not-p, but p.
3. It is possible that X V's, but it is not possible for X to V.
4. My car is in the Main Street parking lot, but it might not be there.
5. If p had been the case q would have been the case, but if p had been the case q might not have been the case.

We have strong reason to believe that each of 1–5 is consistent. The consistency of 1 is addressed in the text, above, and is obvious in any case. The consistency of 2 and 4 are just about as obvious: Surely the mere possibility of not-p is consistent with p. If these were inconsistent, you could infer that not-p from the mere possibility of not-p, which seems absurd. For reasons to think 3 is consistent, see DeRose (1991: 602–3). Following David Lewis (1973: 2, 21–4), many subscribe to the 'duality thesis' of the relation between 'might' and 'would' counterfactual conditionals, and would therefore hold that instances of 5 are genuinely inconsistent. However, the duality account is wrong, and instances of 5 are consistent, as I show in DeRose (1994) and (1999b: 385–413). As I have argued, the hypothesis that one represents oneself as knowing something when one asserts it provides the best treatment of these 'clashing conjunctions'—a treatment that happily does not make any of 1–5 out to be a genuine inconsistency: I address 1–3 at DeRose (1991: 596–605); the treatment of 2 is expanded to cover statements like 4 at DeRose (1998: 72–3); 5 is addressed at DeRose (1999b: 388–90).

generally right—namely, that how well positioned one must be with respect to p to be able to properly assert that p is a variable and highly context-sensitive matter? That's no problem—for the contextualist. The context-variability in what we are positioned to assert is just what the knowledge account of assertion would lead us to expect if what counts as knowledge is a context-variable matter. The contextualist about knowledge who also accepts the knowledge account of assertion welcomes the context-variability in what we can assert, and indeed should have been quite worried if that variability was not found. Following David Lewis's nice phrasing (though Lewis was writing of a connection other than that between assertability and knowledge), the contextualist can say:

I am not one of those philosophers who seek to rest fixed distinctions upon a foundation quite incapable of supporting them. I rather seek to rest an unfixed distinction upon a swaying foundation, claiming that the two sway together rather than independently.[22]

Given contextualism, the knowledge account of assertion naturally takes a relativized form:

The Relativized Knowledge Account of Assertion (*KAA-R*): A speaker, S, is well-enough positioned with respect to p to be able to properly assert that p if and only if S knows that p *according to the standards for knowledge that are in place as S makes her assertion.*

On KAA-R, then, the rule for assertion is: One must assert only what one knows according to the standards for knowledge that are in place as one makes one's assertion. And what one represents as being the case by asserting that p, then, is that one knows that p according to the standards for knowledge that are in place as one makes one's assertion.

Note that KAA-R is one of several possible ways of relativizing KAA. Among the candidates I have here passed over, the most notable is perhaps:

KAA-R2: A subject, S1, is well-enough positioned with respect to p to make true the claim of a speaker, S2, that S1 is 'warranted in asserting that p' if and only if S1 knows that p according to the standards for knowledge that are in place in S2's context as S2 makes her claim.

[22] Lewis (1973: 92). Lewis is there writing about the link between similarity and counterfactual conditionals. Though Lewis was a contextualist (see especially Lewis 1996), he would not have tied assertability to knowledge, as I advocate, because he tended to gravitate toward a probability, rather than a knowledge, account of assertion. However, the probability account Lewis used was well suited to accommodate the context-variability of assertability. Lewis held that one could properly assert something if its probability were 'sufficiently close to 1' (Lewis 1976: 297). It would be easy, and indeed quite natural, to then claim that the matter of just how close is 'sufficiently close' varies with context, which would allow assertability and sufficient probability to 'sway together'.

KAA–R2 would have it that 'S is warranted in asserting p' and 'S knows that p' quite generally 'sway together' if contextualism is true (and stay put together otherwise). But the contextualized version of KAA that I endorse, KAA–R, more appropriately ties the standards a speaker must meet for her assertion to be warranted to the standards for knowledge that govern *her own* context. On KAA–R, then, if contextualism is true, our ascriptions of warranted assertability do not generally sway together with our knowledge attributions; it can happen that 'S is warranted in asserting that p' is true in our context, while 'S knows that p' is false in our context, if the standards for knowledge that are in place in our context are higher than those that govern S's context in such a way that S's belief that p meets the lower standards for knowledge governing her own context, but not the higher standards that govern our context. But on KAA–R, the warranted assertability of p for S *does* sway together with S's own use of first-person knowledge claims—'*I* know that p'—if contextualism is true (and they stay put together otherwise).

Note that the 'swaying together' (or staying put together) posited by KAA–R, though more limited than that posited by KAA–R2, is enough to allow the contextualist to accommodate the grounds for KAA. In particular, note that the swaying together posited by KAA–R allows the contextualist using it to retain the knowledge account of assertion's account for the clash of the Moorean sentences of the form, 'P, but I don't know that p': When a speaker asserts the first half of that troubled conjunction, on KAA–R, she represents herself as knowing that p according to the standards that govern her context as she asserts the conjunction. Thus, when she goes on to assert in the second half of the sentence that she doesn't know that p according to the standards that govern her context, she is involved in a contradiction between one thing that she asserts, and another thing that she represents as being the case. In a similar way, KAA–R retains the knowledge account of assertion's ability to account for the other clashing conjunctions, the handling of which is advanced as a selling point for the account in n. 21, above. KAA–R, of course, could *not* account for a clash produced by the third-person 'P, but *she* does not know that p', but that's to the good, since there is no clash there to account for.

So KAA–R will now become our official formulation of KAA. Note that the invariantist—and, in particular, our target in this chapter, the classical invariantist—should have no objections to letting KAA–R become our official formulation of KAA. The positive grounds we've seen for KAA carry over to KAA–R, and relativizing the account in the way we have begs no questions against the invariantist. KAA–R is relativized to allow for the *possibility* of contextualism, but does not presuppose contextualism. The invariantist should have no special objections to KAA–R that don't equally apply to

an unrelativized version of KAA—except perhaps that KAA-R is needlessly complex. A classical invariantist will simply hold that the last phrase of KAA-R's formulation—'according to the standards for knowledge that are in place as S makes her assertion'—always refers to the same standards, the one-and-only standards that govern all uses of 'know(s)'. KAA-R allows for, but doesn't demand that, the standards that govern the warranted assertability of a speaker's claim that p sway together with the epistemic standards that would provide the relevant truth-condition for her use of 'I know that p', for it equally allows that these two standards stay put together. For the remainder of this chapter, then, we will take KAA-R to be what we mean by the 'knowledge account of assertion' or by 'KAA'.

So, as we've just seen, the contextualist who accepts the knowledge account of assertion will hold that the standards that govern whether an assertion is warranted vary along with the standards for knowledge that govern the asserter's context. What of the classical invariantist who accepts KAA? Such a character is in serious trouble. Given classical invariantism, the knowledge account of assertion is an untenable attempt to rest a madly swaying distinction upon a stubbornly fixed foundation. Less metaphorically, it is an attempt to identify what is obviously a context-variable standard (the standard for the warranted assertion of 'p') with what one claims is a context-invariable standard (the relevant truth-condition of 'S knows that p', according to the classical invariantist). The knowledge account of assertion demands contextually variable epistemic standards for the truth of 'I know that p', and is simply incredible without them.[23]

'But wait! Even the classical invariantist accepts that varying epistemic standards somehow govern knowledge attributions, including, most relevantly here, first-person claims to know. Her only difference with the contextualist is that she holds that it is only the warranted assertability conditions, rather than the truth-conditions, of such attributions that vary with context. So can't the invariantist accept something at least very much like the knowledge account of assertion, but escape trouble by tying the assertability of p to the assertability, rather than the truth, of "I know that p"?'

No. The relevant warranted assertability condition for 'p' cannot be plausibly equated with the relevant warranted assertability condition for 'I know that p'.[24] The connection between knowledge and assertion that works is one

---

[23] Williamson mentions the possibility of contextualizing the knowledge account of assertion (2000: 254–5), but is non-committal about it, writing only that his account 'permits such contextual variation'.

[24] Of course, there are many warranted assertability conditions for assertions. We are here interested only in those that pertain to one's being positioned well enough with respect to the proposition

that ties the relevant warranted assertability condition of 'p' with the relevant *truth*-condition of 'I know that p'. That identification provides a strong argument for contextualism, as we will discuss below in sections 8–9. But alternative potential connections, including that suggested above on behalf of the invariantist, hover about us, and are only subtly mistaken. So before putting the knowledge account of assertion to use, we should look carefully at the connection between knowledge and assertability to make sure we have it right.

7. Assertability and Knowledge: Getting the Connection Right

In section 8, we will follow a fascinating (but largely overlooked) earlier paper by Robert Hambourger (1987) in arguing for contextualism about knowledge from the context-variability of assertability. Unfortunately, Hambourger does not locate a secure path to get him from his solid premiss to his desired conclusion, and instead gets entangled in subtly false claims about the relation between knowledge and assertability. Since these misfiring attempts to connect knowledge with assertability are tempting, it's worth discussing them.

According to Hambourger, the standard, in terms of how well positioned one is with respect to p, that one must meet to be able to properly assert that p is the same as the standard, in terms of position with respect to p, that one must meet to properly claim to know that p[25]—which, in turn, is equated with the standard one must meet for one's claim to know to be true.[26] Both equations of standards—(1) those for properly asserting that p with those for properly asserting that one knows that p, and (2) those for properly asserting

asserted to be able to assert it, and are ignoring other conditions—for example, that the assertion be conversationally relevant.

[25] This is the thesis of part 2 (248–55) of Hambourger's paper. Focusing too much on the notion of evidence, as he does throughout his paper, Hambourger puts that thesis as follows: 'The amount of evidence needed to be epistemically justified in claiming to know a proposition (relative to a given standard of caution) simply is the amount needed to be epistemically justified in asserting the proposition itself straight out (relative to that standard)' (249). I have converted Hambourger's talk about how much evidence one has into the more general terms of how well positioned one is.

[26] This is one of the main conclusions of part 3 (255–62) of Hambourger's paper. He writes, 'I conclude, therefore, that the amount of evidence needed to know a proposition is not fixed once and for all, but that it varies with standards of caution, and, in particular, that the amount needed to know a proposition under any given standard is just that needed to be epistemically justified in claiming to know it…under that standard' (260). Again (see the immediately preceding note), Hambourger expresses these as standards for how much evidence one must have, and I have again converted that to the more general matter of standards for how well positioned one must be.

that one knows that p with those for actually knowing that p—are mistaken, as I trust the considerations below will show to anyone who has deliberated over close calls about whether one is positioned well enough to claim to know that p or should cool one's heels and only assert that p. (We will discuss these matters in a way that is neutral between contextualism and invariantism.)

We are often unsure about whether we know. Sometimes, beyond merely being unsure, we positively go wrong about what we know. One way such a mistake can occur is that we are basing our judgement that we do know on some false belief about some underlying factual matter relevant to the issue of whether we know. The fake barn case again provides an example: Henry quite reasonably thinks that he knows that he's seeing a barn, because he thinks he is in a very normal situation; unbeknownst to him, he's in a region teeming with fake barns and in fact doesn't know. But we can be unsure whether we know even without being mistaken about our factual situation: The factual suppositions on which you're operating are all correct, but your grasp of your factual situation is a bit incomplete, and it's a very close call whether or not you know, as you yourself will say if asked whether you know. You just can't tell whether you know. Well, then, you're in no position to assert that you do know. But on some of those occasions, of course, you really do know: Are we to suppose that all such close calls between knowing and not knowing occur in situations where in fact one does not know? (If so, then realizing that all close calls are actually cases of non-knowing would make it all too easy to correctly handle a close call!) But if you know, you meet the epistemic standards for actually knowing. And if, as we've already determined, you are not positioned to assert that you know, then you don't meet the standards for asserting that you know. Hence, equation (2) is invalidated.

There is a certain kind of counter-example to equation (2) that I can't give here: I can't describe in more concrete terms a situation where a subject clearly knows but clearly is in no position to assert that she knows. I had to describe the situation as involving close calls. But if the calls are close, they won't be clear. Nevertheless, we can see by the above argument that there must be some unclear cases where one knows but is not yet well enough positioned to assert that one knows.

Such unclear scenarios also give the lie to equation (1), but let us proceed in more picturesque terms this time. Imagine a situation in which you're clearly positioned well enough with respect to p to properly assert both that p and that you know that p: You see a cat just a few feet in front of you, the situation is favorable, and you are clearly able to properly assert that you are seeing a cat and that you know that you're seeing a cat. Then consider a series of cases just like the one just described, except that in each case you

are a bit further from, and therefore get a little worse look at, the cat. At the end of the series, you are clearly in no position to assert even that the very distant thing you're seeing is a cat, much less that you know it to be a cat. Toward the middle of this series of cases, you get many close calls, both for whether you're able to assert that p and for whether you can properly claim to know that p. I don't think we will find any particular distance at which we get a clear counter-example to equation (1)—a distance at which it's clear that you can properly assert that p but clearly cannot properly claim to know that p. Nevertheless, as we can all sense, throughout the intermediate cases, the knowledge claim is more problematic than is the simple assertion that p, and it seems clear that tougher standards govern the propriety of the claim to know that p than govern the simple claim that p, even if the standards are too close to one another to yield any cases where a speaker clearly meets one but clearly falls short of the other. (Compare: I'm grading philosophy essays on a 100-point scale. My standards are clearly tougher for an essay meriting a score of at least 77 than for one earning a score of at least 75, despite the fact that there won't be any essays that to me both *clearly* meet those lower standards and *clearly* fail to meet the higher.)

It's often tough to properly assert. Even where you have a reasonable belief that p, it's far from automatic that you're in any position to flat-out assert it. And if you feel that the issue of whether you know that p looms large in deliberating over whether you are in a position to assert that p, then, of course, being an advocate of the knowledge account of assertion, I sympathize with that feeling. Nevertheless, as I hope we can all sense, going beyond the simple assertion of p to the claim that one knows that p is, as Austin put it, 'taking a new plunge' (1979a: 99), and, contrary to equation (1), is evidently subject to higher epistemic standards.

What *is* true is the equation that (1) and (2) together imply: that of the standards for knowing that p with those for being in a position to properly assert that p. That equation is the knowledge account of assertion: One must know that p to be positioned well enough to properly assert that p.

While, like equations (1) and (2), this equation maintains an important tie between knowledge and proper assertion, it also can explain why claiming to know is to take a greater 'plunge' than simply asserting p, why it's hard to come up with clear counter-examples to equations (1) and (2) despite those equations' evident falsehood, and can perhaps even help to explain the appeal of the KK principle (the claim that if S knows that p, then S knows that S knows that p).

Simple as it is, I've found that Table 3.1 nevertheless helps in making the needed points.

Table 3.1

In asserting	S (of course) asserts that	And S represents it as being the case that
'p'	p	S knows that p
'I know that p'	S knows that p	S knows that S knows that p

Table 3.2

S's assertion	Truth-condition of assertion	Warranted assertability condition of assertion
'p'	p	S knows that p
'I know that p'	S knows that p	S knows that S knows that p

Alternatively, we can relabel our columns in terms of the truth-conditions and the relevant warranted assertability condition (see Table 3.2).

As Table 3.1 shows, the knowledge account of assertion does equate what S asserts with 'I know that p' with what S represents as being the case by asserting 'p', and, as Table 3.2 shows, it equates the relevant warranted assertability condition of simple assertion with the truth-condition of the knowledge claim. However, despite these equations, when you compare apples with apples and oranges with oranges, the knowledge claim (on the bottom of each table) is stronger than the simple assertion (directly above it), both with respect to what you assert (the middle column) and with respect to what you represent as being the case by making the assertion and the relevant warranted assertability condition of your assertion (the right column). Thus, the account respects the fact that making the knowledge claim really is taking more of a plunge than is simply asserting that p.

Why are counter-examples to the mistaken equations tough to come by? Equation (1) mis-equates the warranted assertability conditions for 'p' with the warranted assertability conditions for 'I know that p'. Equation (2) mis-equates the truth-conditions of the knowledge claim with the knowledge claim's warranted assertability conditions. In each case, according to the knowledge account of assertion (see Table 3.2), this is a misfiring attempt to equate something stronger (S's knowing that S knows that p) with something slightly weaker (S's knowing that p). But, given that it's difficult to come up with a case where a subject clearly knows some p, while clearly failing to know that she knows that p, the knowledge account of assertion can then explain why it's difficult to construct clear counter-examples to equations (1) and (2).

Also, according to the knowledge account of assertion, one represents it as being the case that one knows that one knows that p when one claims to know that p. Given a general tendency to confuse what a speaker merely represents as being the case in making an assertion with what the speaker actually asserts, this feature of the account, together with the already noted difficulty in coming up with clear cases where a subject knows without knowing that she knows, may go a long way toward explaining the attraction that the 'KK principle' holds for some.[27] (Here I must tread lightly, since the principle has never held much allure for me.)

8. The Argument from Variable Assertability Conditions

The knowledge account of assertion provides a powerful argument for the contextualist to wield against the classical invariantist. And supposing subject-sensitive invariantism can be shown to be wrong, this constitutes a powerful argument in favor of contextualism. Recall that the classical invariantist denies that varying epistemic standards provide the truth-conditions for knowledge attributions, including first-person claims to 'know'—though he will typically claim that varying epistemic standards provide the warranted assertability conditions for knowledge attributions. But if, as KAA would have it, the standards for when one is in a position to warrantedly assert that p are the same as those that constitute a truth-condition for 'I know that p', then if the former vary with context, so do the latter. In short: KAA—which, recall, we're construing as KAA-R—together with the context-sensitivity of assertability[28] yields the conclusion that the *truth-conditions* of 'I know that p' are in fact context-sensitive in the way the classical invariantist denies.

One thing that one looks for in an argument is that its premiss be more secure than the conclusion it is being used to establish. This points to a question: If KAA is right to equate the standards for the warranted

---

[27] What might also help in explaining the attraction of the KK principle is locating a principle that is in the vicinity of KK but is weaker and more defensible. The principle that knowing is a 'weakly luminous' condition, in the sense I describe in DeRose (2002b: 576–7), may serve this purpose. Briefly, to say that knowing is weakly luminous, in this sense, is to say that whenever one safely knows (whenever one knows and is not close to not knowing), one is in a position to know that one knows.

[28] Again, I am speaking here of a particular kind of warranted assertability: that having to do with whether one is in a position to make the assertion one is making. I take it as obvious that there is context-variability in this aspect of the conditions for warranted assertability.

assertability of 'p' with those for the truth of 'I know that p', then why is the claim that the former are context-variable any more secure than the claim that the latter are? What makes the first claim a suitable premiss for the second?

It is clear and uncontroversial that knowledge attributions, including first-person claims to know, are *in some way* governed by varying epistemic standards. What make contextualism controversial are warranted assertability concerns—the worry that these varying standards are conditions only for warranted assertability, rather than the truth, of the attributions. This worry about our conclusion does not apply to our premiss, which itself already concerns warranted assertability.

It is difficult to deny that the matter of how well positioned one must be with respect to a matter to be able to assert it varies with context: What one can flat-out assert in some 'easy' contexts can be put forward in only a hedged manner ('I think…', 'I believe…', 'Probably…', etc.) when more stringent standards hold sway. Even classical invariantists, who deny that the truth-conditions of knowledge attributions are sensitive to varying standards, tend to agree that the *warranted assertability* conditions of knowledge attributions vary with context. (Indeed, as we've seen, they attempt to use that fact in explaining away the allure of contextualism.) And it's clear that this is true not only of knowledge attributions, but of assertions generally—clear enough that I trust that illustrative examples are not needed here. And no 'Warranted Assertability Maneuver' can be wielded against us here, for the strong intuitions we are utilizing are themselves judgements concerning warranted assertability.

Given how secure is the premiss that assertability is context-variable, KAA, which provides the only other premiss needed by our argument, is lethal to classical invariantism. It turns against the classical invariantists the very context-sensitivity in warranted assertability that they themselves habitually appeal to. Unfortunately for classical invariantists, the knowledge account of assertion is not only fatal to them, but also very well supported.

9. An Argument for Contextualism?

The argument we've been considering shows, against the classical invariantist, the conclusion that the truth-conditions of first-person knowledge claims vary in the way the contextualist claims they do. But does this argument establish contextualism?

One quite strange possible position that can accept the conclusion mentioned above while stopping short of generally accepting contextualism would have it that varying standards sensitive to the subject's context govern whether that subject can truthfully claim to 'know' using a first-person knowledge claim, while a single, invariant standard, not sensitive to anyone's context, governs third-person attributions. (We might as well leave open what such a strange view will do with second-person attributions.) But nobody—at least that I know of—endorses this ugly and unmotivated amalgamation of contextualism and classical invariantism.

But, as Thomas Blackson has pointed out (2004), our argument from assertability also fails to rule out subject-sensitive invariantism.[29] SSI also agrees with contextualism that varying epistemic standards govern whether a speaker can truthfully claim 'I know that p', while denying contextualism. But SSI avoids the ugliness of the amalgamation we considered above, for SSI holds that varying standards also govern third-person knowledge attributions. For on SSI, features of a subject's context determine a standard that governs whether the subject herself, or any other speaker, including those not engaged in conversation with the subject, can truthfully say that that subject 'knows'. On SSI, then, the standards governing whether (first-, second-, and third-person) sentences attributing knowledge to a subject are true do vary with the subject's context: If the subject had been in a different situation, different standards would have provided the truth-condition for those sentences. So the defender of SSI can accept the conclusion that truth-conditions of first-person knowledge claims vary with the speaker's context, for SSI holds that the standards vary with the subject's context, and in the case of first-person knowledge claims, the speaker *is* the subject. But, as we discussed in section 11 of Chapter 1, SSI is not a contextualist view, and does not, for instance, allow what contextualists insist on: that one speaker can truthfully say that a subject 'knows', while another speaker, in a different and more demanding context, can say that the subject does 'not know', even though the two speakers are speaking about the same subject 'knowing'/'not knowing' the same proposition at the same time.

I have long thought that how we use third-person knowledge attributions shows that views like SSI cannot be right,[30] and have therefore long thought

---

[29] Not worrying about views like SSI, in DeRose (2002a) I presented the above argument as an argument for contextualism, and Blackson rightly objected that SSI escapes the argument.

[30] I've long played around with third-person cases like those I use against SSI in Chapter 6, but, as nobody seemed to take such a view seriously before the relevant work of Hawthorne, Fantl and McGrath, and Stanley, I only published very brief discussions of such cases—see DeRose (1996b: 112; 1999a: 191)—before those discussions of SSI came out.

of classical invariantism as the real threat to contextualism. Thus, when I formulated this argument that favors contextualism over what I construed as its only threatening rival, I construed it as an argument for contextualism. But since SSI is now a serious contender, this argument from assertability constitutes a case for contextualism only if SSI can be ruled out—as I will attempt to do, especially in Chapter 7.

10. The Generality Objection Defeated

Having discussed the Knowledge Account of Assertion and the argument for contextualism that it provides, we will now, as promised, continue our campaign against classical invariantism by utilizing KAA against the Generality Objection.

Recall that according to the Generality Objection, the fact that knowledge attributions tend to become unassertable as we move to contexts governed by higher epistemic standards is just one part of a very general phenomenon. In such contexts, it becomes more difficult to assert *anything*. There's nothing special about knowledge attributions here. And since this rise in the conditions for warranted assertability does not generally reflect any change in truth-conditions, there's no reason to be found here for thinking that the truth-conditions of knowledge attributions vary with context in the relevant way.

KAA yields an answer to the Generality Objection by providing the contextualist with an alternative and, as it turns out, superior explanation for the harmony the Generality Objector points out: Since, as KAA dictates, the relevant warranted assertability condition for the assertion of the simple 'p' is that the speaker know that p according to the epistemic standards that govern her assertion, as those standards for knowledge go up, making 'I know that p' go false and unassertable, the standards for the warranted assertability of the simple 'p' will rise, making it unassertable too.[31] Thus, it's no surprise

---

[31] In the text here, I explain why the assertability of 'p' fades away together with that of first-person claims to 'know' that p, where the phenomenon being accounted for is relatively straightforward. Recall from our discussion in section 4, however, that there is also something to account for in the case of third-person knowledge attributions. As I noted in section 4, it often happens that a speaker is in a position to assert 'p' but is clearly in no position to assert that a certain S 'knows' that p, sometimes simply because the speaker doesn't know anything about S. Still, if we start with a third-person case in which 'p' and 'S know[s] that p' actually are both warranted, and then raise the standards in our imagination, we find that the two claims move toward unassertability together. That's because 'S knows that p' requires for its truth the truth of p. Thus, to 'know' according to the speaker's own standards that S knows that p, and thus to be in a position to assert 'S knows that p', the speaker must know according her own current standards that p is true. And, on our new account based on KAA-R, that

that the assertability of the knowledge claim and that of the simple claim fade away together as we move to contexts governed by higher standards.[32] But on our new account (based on contextualism and the knowledge account of assertion), there *is* something special about knowledge here. The variability in the epistemic standards governing knowledge is *not* just one more aspect of a general variability in the epistemic standards for assertion, and certainly does not derive from the more general variability. Rather, as the standards go up, it is *because* less and less counts as knowledge that fewer and fewer things become assertable. The general variability in the standards for warranted assertability is explained by the variability in the standards for knowledge. And, according to our new account, when the epistemic standards go up past the speaker's ability to meet them, 'I know that p' not only becomes unassertable, as the simple 'p' does: It becomes false.

This suggests a test for deciding between the two accounts on hand: If, as our new account would have it, 'I know that p' becomes false, and not just unassertable, as the epistemic standards rise above the speaker's ability to meet them, then speakers should not just become reluctant to claim to know, but should go so far as to deny that they know. Consideration of the relevant conversational phenomena reveals that our new account is superior to the Generality Objector's account on just that score.

First, consider what happens to a simple 'p' when we move from a context in which one counts as being in a position to assert that p to a context in which one does not so count. Not positioned well enough to assert that p, what can one say instead? Often, some form of hedged claim is a good choice in such a situation: Instead of flat-out asserting p, one can, for instance, say, 'Probably p', 'I'm pretty sure that p', 'I think that p', etc. Or perhaps one can assert some weaker proposition less specific and less ambitious than p that one is still positioned well enough with respect to, despite the elevated standards. And in some situations silence is a good choice. But one thing one cannot do is to assert that not-p! In the 'high-standards' Bank Case, for example, the epistemic standards seem too elevated for me to properly declare flat-out, 'The bank will be open on Saturday.' But, of course, neither am I in any position to

the speaker knows that p according to the standards currently governing her talk is precisely what, on the knowledge account of assertion (KAA-R), the speaker represents as being the case when she simply asserts that p.

[32] To the extent that the assertability of the knowledge claim and the simple claim don't fade away exactly together because the knowledge claim is somewhat more problematic in some intermediate contexts, our account, by making the warranted assertability condition for the knowledge claim a bit more stringent than that of the simple claim (see especially our discussion of the tables in section 7), does a good job of handling this.

say, 'The bank will not be open on Saturday'! I am not in a position to say *that* even according to the more relaxed standards governing the 'low-standards' Bank Case, and, unsurprisingly, the rise in standards does not make it easier to make the assertion.

If the Generality Objector were correct that the context-sensitivity displayed by knowledge attributions is just more of the same, we would expect similar behavior from 'I know that p'. But in fact knowledge claims behave very differently. In many cases, when we move into more demanding contexts where 'I know that p' becomes clearly unassertable, the speaker will do something the Generality Objector should find shocking: She will go so far as to deny that she knows, admitting, 'I don't know that p'.

Indeed, the pairs of cases contextualists have typically used to motivate their view have displayed just such verbal behavior: 'Low-standards' cases, where an attribution of knowledge seems perfectly in order (is just what a normal speaker would/could say) are paired with 'high-standards' cases in which speakers, describing similarly positioned subjects, not only seem prohibited from properly describing similarly positioned subjects as knowers, but, with complete propriety, go so far as to describe those subjects (often themselves) as non-knowers.[33] As I've already noted, defenders of invariantism have wisely accepted that these cases reflect how we use 'knows'. They have merely attempted to claim that this behavior can be explained just as well without the contextualist's appeal to varying truth-conditions for knowledge attributions.

But the Generality Objection seems able to handle only our reluctance to claim knowledge, and seems ill suited to explain why we go so far as to deny that we know. The Generality Objection would have us expect the former, since it predicts that, like all other assertions, claims of knowledge will become harder to make as the epistemic standards go up. But it is at a complete loss to explain the evident fact that 'I do *not* know' often becomes assertable as the standards go up, when it was unassertable at lower standards.

The invariantist may claim that when 'I know that p' becomes unassertable, we mistake its unassertability for falsehood, and, overreacting to the claim's unassertability, we go overboard and go so far as to assert its negation. But, as I've already pointed out, we do not do this with the simple 'p'. And it's no close call, either: It seems *crazy* to assert 'not-p' in cases where we are not in a position to assert that p due to elevated standards, though we would meet lower standards for being positioned to assert that p. The Generality Objection

---

[33] This is true, for instance, of both Cohen's Airport Case (see Chapter 1 n. 1) and my Bank Cases (described in Chapter 1, section 1).

leaves it quite a surprise that when it comes to the knowledge claim, we do go so far as to deny it.

But our alternative explanation, based on contextualism and the knowledge account of assertion, handles all this nicely. It too explains why both 'p' and 'I know that p' become unassertable as the standards come to clearly exceed those the speaker can measure up to. But it is also no surprise on our alternative account, based on the knowledge account of assertion, that as the standards go up to a level at which the subject's belief that p clearly does not meet them, she will be described as not knowing that p.

11. Check the Negations!

Where it is questionable whether a claim in a situation is or would be false, or, on the other hand, merely inappropriate, a good piece of general methodological advice (like the general advice to 'Follow the money' in a certain type of criminal investigation) is: Check the negations (of the relevant claims).[34]

As we've seen, following that advice shows us where the Generality Objection stumbles: While 'I know that p' and 'p' behave similarly enough to make the Objection look promising, with both becoming unassertable, and roughly together, as the epistemic standards go up, the glaring difference

---

[34] In checking the negations of the claims in question, one must keep in mind that not all statements of the form 'not-p' are genuine denials of p. There is a special type of case where speakers will say 'not-p', where they are merely denying the assertability of p. The classic treatment of this general phenomenon, known as 'metalinguistic negation', is Horn (1989). Such cases are distinguishable from genuine denials of p in that cases of metalinguistic negation involve an unusual 'fall–rise' intonation and a follow-up: they are followed by what should instead be said. A special case of this that is quite common is where speakers (at least appear to) deny, with the unusual intonation, something weaker and then go on to say something stronger: 'He's not *big*, he's *huge*'; 'She's not *smart*; she's *brilliant*.' Here, it seems, what's being conveyed is that, for instance, he isn't *only* big, and since something stronger—'He's huge'—should instead be said, the weaker thing—'He's big'—should not be asserted. I take it then that these should not be understood as really being denials of the weaker claims ('he's big', 'she's smart') or as assertions of the negations of the weaker claims, but rather merely as denials that the weaker claim should be made. Where a flat-footed (unemphasized and without unusual intonation) and unfollowed denial seems wrong, but the same 'denial' can appropriately be said with the right emphasis and follow-up, we should conclude that it is really only the assertability, rather than the truth, of what appears to be denied, that is really being denied in the special cases (of emphasis and follow-up). And if an assertion of 'not-p', made without the special intonation and follow-up, seems correct, we are justified in taking it to be a genuine denial of p, and not merely a case of metalinguistic negation, for metalinguistic negation seems to require the intonation and follow-up. To see this, consider a case of genuine metalinguistic negation—say, 'He's not *big*, he's *huge*.' In such a case, it would be clearly wrong for the speaker to simply and flat-footedly (without the special intonation) state, 'He's not big.'

between 'I don't know that p' (which often becomes assertable when the standards go up) and 'not-p' (which doesn't) spells the doom of the Generality Objection, especially given the presence of an alternative account of the relevant phenomena that outperforms it here.

Our willingness to admit in 'high-standards' contexts that we don't know things we do claim to know in less demanding contexts also casts other imposing shadows over the prospects that any warranted assertability maneuver will work for the invariantist here. Beyond wiping out the Generality Objection, which is the only form of such a maneuver (that I know of, at least) that looked at all promising, it also presents a formidable general obstacle to the success of WAMs in this case, because this aspect of the behavior of 'I know'/'I don't know' puts any such a WAM on the failing side of further criteria (beyond the generality criterion we detected above in section 3) that we can discern for successful WAMs.

Consider first a reasonable WAM—the one we considered earlier, in defense of the DKO account of 'It's possible that p_{ind}', will do. Recall that the problem cases for DKO are those where the speaker knows that p is true. About such cases, many have some inclination to think that it would be false for the speaker to say, 'It's possible that p_{ind}' (though, again, many will instead or in addition find themselves with some opposing inclination to think the claim would be true, though inappropriate and misleading). We have already seen how the defender of DKO can use the very general 'Assert the Stronger' rule to disarm or at least mitigate the force of this intuition that is hostile to his theory. Auxiliary considerations provide further leverage for this maneuver. First, of course, is the fact that the hostile intuition (that 'It's possible that p_{ind}' is false in the relevant situation) is for many not very strong and/or is balanced by a conflicting intuition.

And if we 'Check the Negation', we find two further good-making features motivating the WAMing of this intuition. First, where a speaker knows that p, while it does indeed seem wrong, and may even seem to some false, for her to assert 'It's possible that p_{ind}', it seems at least as bad—in fact, considerably worse—and *certainly* seems false, for her to say instead 'It's *im*possible that p_{ind}' or 'It's *not* possible that p_{ind}'.[35] But it seems quite unlikely that *both* 'It's

---

[35] But recalling the previous note, it seems that 'It's not possible that p' can be said where the speaker knows that p, under the special circumstances that the 'denial' of the possibility is said with the right 'fall–rise' intonation, and is immediately followed by a stronger statement: e.g. 'It's not *possible* that the book is there; it's *certain* that it is' or 'It's not *possible* that the book is there; it *is* there'. Since such cases of intonation-and-follow-up fit the general pattern mentioned in the preceding note, I conclude that they would not after all be real denials of the possibility (note, too, that 'It's impossible' still seems wrong), but are just ways of saying that the mere possibility is not what should be asserted. My argument in

possible that p_{ind}' and 'It's *not* possible that p_{ind}' are false, so we have good reason to believe that something is not as it seems here; we're going to have to explain away the misleading appearance of falsehood of *something*. By contrast, in the problem cases for the crazed theory of bachelor (cases involving married males), or for Fraction's theory of knowledge (cases involving false beliefs), there's no such pressure to have to explain away any appearances: It seems false to say of married men that they are bachelors, and it seems true to say of them that they are not bachelors; it seems false to describe false beliefs as cases of knowledge, and it seems true to say that believers of falsehoods do not know what they believe. So there appears to be no problem in these cases, making it all the more problematic to claim that in fact there is a *big* problem here, with the intuitions of both falsehood and truth being mistaken.

Second, and closely related, having to explain away intuitions of truth as well as of falsehood not only increases the *number* of intuitions that our lame maneuvers must explain away, but also places on them an explanatory burden of a different and more problematic *kind*.

As we have seen, the plausible WAM used in defense of DKO appeals to the generation of a *false implicature* to explain away the appearance of falsehood in the problem cases for the theory. It is not surprising that a true assertion will be inappropriate, and may seem false, if it generates a false implicature: We should avoid conveying falsehoods, whether the falsehood conveyed is part of truth-conditional content of the assertion or is an implicature generated by the assertion, and where *something* false *is* conveyed by the making of an assertion, it's not surprising that we might mistake that for the assertion's itself being false. By contrast, it seems much more problematic to claim that an assertion that seems true is in fact false by means of a claim that, though the assertion itself is false, it generates a true implicature and is therefore a warranted assertion. For, except where we engage in special practices of misdirection, like irony or hyperbole, don't we want to avoid falsehood both in what we implicate and (especially!) in what we actually say? So, it would seem that a false assertion will remain unwarranted, despite whatever true implicatures it might generate.[36] Thus, attempts to explain away appearances of truth by appeal to the generation of true implicatures seem more problematic than attempts to explain away appearances of falsehood by appeal to the generation of false implicatures. It

---

the text depends on the fact that in situations where a speaker knows that p, it seems clearly wrong for her to simply and flat-footedly state, 'It's not possible that p'. That she could appropriately state those same words with special intonation and follow-up is irrelevant, for then she would be engaging in the special practice of metalinguistic negation.

[36] Note the connection here with the aspect of my methodology explained in section 2–3 of Chapter 2.

is then no wonder that most clearly effective WAMs are of the latter type. Indeed, none I know of involve explaining away apparent truth by appeal to the generation of true implicatures.[37]

Putting all our criteria together, we can see how much our proposed WAM in defense of the DKO account of 'It's possible that p_{ind}' has going for it. It starts out with the advantage of being aimed at an intuition that is a good candidate for being WAMed: an intuition that an assertion ('It's possible that p_{ind}' in cases where the speaker knows that p) would be false that is opposed by an intuition that the opposite assertion would also be false in the relevant circumstances. It then explains away the relevant apparent falsehood by appeal to the fact that a false implicature would be generated by the assertion in the relevant situations—which, as we have seen, is less problematic than is trying to explain away the appearance of truth by appeal to the generation of a true implicature. And finally, it explains how the false implicature in question gets generated by appeal to very general rules, rather than to ad hoc rules that attach only to sentences of the type in question.[38] By contrast, Fraction's defense of the 'Equivalence Thesis' and our imagined defense of the crazed

---

[37] I am limiting my attention here to effective WAMs that are aimed at what might otherwise really be tempting treatments of sentence types. I take it that the defense of DKO we are considering falls into this camp: One might really be tempted by the DKEW account of 'It's possible that p_{ind}'. Likewise, to give just one well-known example of a clearly successful WAM, initially one might be tempted to think that it is not true to say of a bright red (British) mailbox that is only two feet away from one in bright sunlight 'It looks red to me', and one might thereby seek an account of 'looks red' that respects this appearance (Grice 1989: esp. 227–30). However, I take it that Grice has successfully WAMed the intuition that 'looks red' cannot be truthfully applied in such a case. Note that this WAM gets on the good side of all my tests: It seeks to explain away an appearance of non-truth; it does so by means of a general conversational rule; and the appearance of non-truth that it seeks to explain away is balanced by a similar appearance of non-truth that attaches to the opposite statement: In the circumstances under discussion, it also seems wrong to say that the mailbox 'does not look red'—except, to recall nn. 34 and 35, in the special case where one says this with the proper emphasis and a proper follow-up: 'It doesn't *look* red; it *is* red!' In that case, we are reasonable to understand the claim as a case of metalinguistic negation, rather than as a genuine denial of the proposition that the mailbox looks red.

[38] Patrick Rysiew argues that even the 'doesn't know otherwise' condition (the condition that the speaker not know that not-p) is merely something conveyed by an assertion of 'It's possible that p_{ind}', and is not part of the truth-conditional content of the modal statement. After noting that I accept that 'I don't know that p' is merely conveyed by, and is not part of the truth-conditional content of, 'It's possible that p_{ind}', he writes: 'Curiously, however, [DeRose] seems not to see that an exactly parallel argument establishes that the speaker's not knowing that *not-p* isn't part of the truth-conditions of such statements either' (Rysiew 2001: 510). However, so far from being 'exactly parallel', the arguments are *very* far apart indeed. My case for WAMing the condition I reject, both here and in the paper Rysiew addresses (DeRose 1999a), is based on three considerations I put forward that I argue make this WAM credible. *At least* two of those three considerations *clearly* fail to apply to the 'parallel' case Rysiew proposes (and the remaining one certainly *appears* to fail to apply): (1) Where the condition I reject as a truth-condition (that the speaker not know that p) fails to hold (where the speaker does know that p), while it can seem false for her to say 'It's possible that p_{ind}', it also seems false for her to make the opposite assertion (that p is 'not possible' or is 'impossible'). This makes this condition suspect, and a good candidate for WAMing. The same cannot be said for the condition I accept—that the speaker

theory regarding 'bachelor' falls on the bad side of all the distinctions we have been using.

Unfortunately for classical invariantists, it looks like their attempts to handle the conversational patterns that seem to support contextualism do no better by our reasonable criteria than the desperate and lame maneuvers we have been imagining. As has emerged as we checked the relevant negations, invariantists do not begin with a good candidate for WAMing, and they have to explain away as misleading intuitions of truth as well as intuitions of falsehood. For in the 'low-standards' contexts, it seems appropriate and it seems true to say that certain subjects know and it would seem wrong and false to deny that they know, while in the 'high-standards' context, it seems appropriate and true to say that similarly situated subjects don't know and it seems inappropriate and false to say they do know.[39] Thus, whichever set of appearances the

not know that not-p. Where *it* fails (where the speaker *does* know that not-p), it seems both false for the speaker to assert 'It's possible that p_{ind}' and true for her to deny that p is possible. (2) Thus, in WAMing the 'suspect' condition, I only have to explain away intuitions of falsehood, while Rysiew will have to do not only that, but also explain away intuitions of truth, which, for the reason I've given, seems more problematic. (3) I have a hard time evaluating whether any maneuver of the type Rysiew proposes will be able to successfully generate the result he needs based on his account of the truth-conditions of 'It's possible that p_{ind}' together with general conversational rules, and thus get on the 'good side' of my third consideration, for I can't see what he thinks the truth-conditions of 'It's possible that p_{ind}' are. But no attempt to explain away the condition I endorse can be even close to 'exactly parallel' to the account with which I battle the condition I reject, for, as we have seen, my account relies on the claim that 'p' and 'I know that p' are stronger than (imply but are not implied by) 'It's possible that p'. A closely parallel account targeting the condition I endorse would then have to rely on the very implausible claim that 'not-p' or 'I know that not-p' implies 'It's possible that p_{ind}'. Also, especially in 'Epistemic Possibilities', I use the account of 'It's possible that p_{ind}' that I accept—which includes as a truth-condition what Rysiew rejects as such—to explain *many* aspects of the behavior of the sentence type, and I cannot see how all this can be explained if this account is rejected.

[39] It is crucial to keep in mind that when we speak of 'high-standards contexts' here, we are not thinking of contexts in which the hyperbolic doubts of philosophical skepticism are in play. There the relevant intuitions are not that strong. We are rather thinking of cases like my 'high-standards' Bank Case—cases where doubts are in play that are a bit more remote than are usually taken seriously, but not nearly so remote as the playthings of philosophers, and where it is reasonable for such unusually remote doubts to be taken seriously, given the practical concerns of the situation. About such cases, I submit, the relevant intuitions really are quite strong—and, for reasons given in Chapter 2, are of the type we ought to treat quite seriously.

Rysiew argues that the invariantist attempt to explain the data being used to support contextualism actually satisfies the desideratum of being aimed at what is, in this way, a good target for WAMing (Rysiew 2005: 59–60), but his argument is based on observations about what one can say about whether one knows that one is not a bodiless brain-in-a-vat (a BIV):

But, on the face of it, the sophisticated invariantist view outlined above seems to satisfy DeRose's criteria for a successful WAM. Thus, while it can seem wrong to claim that I know that I'm not a BIV, for example, it can seem no less wrong to say that I *don't* know that I'm not a BIV. (2005: 59)

Now, I certainly admit that it's intuitively quite unclear whether one knows that one is not a BIV, or, more generally, whether one knows that various skeptical hypotheses are false. (For my attempt to explain the various wavering intuitions involved, see esp. DeRose 2004b: 38–9.) And I believe that in

classical invariantist seeks to discredit—whether she says we're mistaken about the 'high' or the 'low' contexts—she will have to explain away both an appearance of falsity and (much more problematically) an appearance of truth. And, as we have seen, the warranted assertability maneuver, as it has been put forward by its advocates, fails our test of generality; and the Generality Objection, which is our attempt to help out the invariantists on this score, while initially promising, ultimately fizzles when we 'check the negations' and discover that the explanation on which the Generality Objection is based is inferior to a rival contextualist explanation.

It is very easy for defenders of invariantism to 'chalk it all up to pragmatics'—to claim that the varying epistemic standards that govern our use of 'knows' are only varying standards for when it would be appropriate to say we know. But it turns out to be no easy matter to make good on such a claim, and we now have little reason to think such claims credible. Not only does the invariantist's use of such a maneuver fail all the tests we have discerned, but, as we saw above in sections 8–9, what a close look at the conditions for warranted assertability actually leads to is a powerful argument *in favor of* contextualism over its chief rival.

Appendix: Rysiew's and Unger's Invariantist Accounts

Toward the end of section 3 of this chapter, I claimed that defenders of classical invariantism have not proposed any credible account on which general rules of conversation together with their proposed invariantist truth-conditions for knowledge attributions predict the pattern of varying conditions of warranted assertability of knowledge attributions and denials that we encounter in natural language. Here, as promised, I discuss what can appear to be a couple of exceptions to that claim: the invariantist accounts of Patrick Rysiew (2001) (further explained in Rysiew 2005) and of Peter Unger (1975).

Our main focus, then, will be on how these invariantist attempts to WAM the conversational data supportive of contextualism fare with respect to our first desideratum for potentially successfully WAMs—that of generating their explanations by the use of

---

many cases, Rysiew's observation above is correct: It seems at least somehow wrong both to say that one knows and to say that one doesn't know that one is not a BIV. But the argument for contextualism under discussion is not at all based on claims about what it's right to say about whether one knows skeptical hypotheses to be false! It is instead based on the apparent truth and appropriateness of claims to the effect that subjects know more ordinary propositions (like that the bank is open on Saturdays) in LOW cases, paired with the apparent truth and appropriateness of denials that subjects know more ordinary propositions in HIGH cases. It's the truth of the most compelling of *those* claims and denials that needs to be WAMed if the argument from ordinary usage for contextualism is to be dodged.

general rules (in addition to accounts of the truth-conditions of the relevant sentences), rather than by positing special rules for the use of the terms in question. Before getting to that, however, it's important to note that, like any invariantist attempt to WAM away the data supportive of contextualism, these attempts will fail to satisfy our other two desiderata that we discerned in section 11. Thus, attempts to WAM the data supportive of contextualism start out with two strikes against them. In these two respects, they are like our examples of clearly ineffective WAMs—like Fraction's lame defense of the 'Equivalence Thesis'—rather than resembling the more plausible WAM we tried in defense of the DKO account of epistemic modal statements. But how do Rysiew's and Unger's attempts fare by our first desideratum, that of generality?

Rysiew

Responding to my challenge to the invariantist to work out an account of the relevant behavior of knowledge attributions—that they are *somehow* subject to varying epistemic standards—using general rules of conversation together with their proposed invariantist account of the semantics of the sentences, Patrick Rysiew (2001) has stepped up to the plate and attempted to do just that.[40] As Rysiew points out (2001: 491), his moderate (i.e. non-skeptical) classical invariantist account is 'underwritten' by the very general conversational rule 'Be relevant!' His account can therefore appear to be of the type of account I have urged us to seek. But, at least as I was initially inclined to understand his account, in addition to there being worries about whether the account can actually explain what is needed, it isn't in the end even the type of account I've been urging us to seek. However, Rysiew has since (2005) made clear that he intended his account to be understood in a different way, on which it really is of the right type. But I believe that understanding it in this way exacerbates the other problems the account has in working at all. Since both construals are worth discussing, we will now look at Rysiew's account, both in the way I was initially inclined to understand it, and also on its correct, intended reading.

As I was initially inclined to understand it, Rysiew's theory of knowledge attributions starts by positing two separate but related meanings that an assertion of 'S knows that p' carries: Such a knowledge attribution semantically imparts that S has a true belief that p and can rule out all the *relevant* alternatives to p, where the standard for relevance is invariant (2001: 488), but also conveys the second, pragmatic meaning that S has a true belief that p and can rule out all of the *salient* alternatives to p, where the salience of alternatives is a highly context-sensitive matter. Where the contextually salient alternatives diverge significantly from the relevant alternatives, the injunction to 'be relevant' gets the listener to fasten on the second, pragmatic meaning, since the issue of whether the subject satisfies the pragmatic meaning of the attribution is more relevant to the conversational purposes at hand than is the

---

[40] Rysiew is responding to the challenge as I made it in (DeRose 1999a); see esp. Rysiew (2001: 491).

issue of whether he satisfies what is semantically imparted by the sentence.[41] Thus, on Rysiew's treatment of my 'high-standards' Bank Case (2001: 490), my claim that I 'don't know' that the bank is open on Saturday is false, but all is well, for my wife can be counted on to ignore the falsehood I strictly said, since mere 'knowing'/'not knowing' is not what is most relevant to our there unusually heightened concerns, and she can be trusted to instead focus on the second, pragmatic meaning that my assertion carries—the truth that it's not the case that I have a true belief and can rule out all of the *salient* alternatives to p, where in such a high-stakes case like this, the set of salient alternatives is quite a bit broader than is the set of relevant alternatives. On this understanding of it, then, Rysiew's account begins by positing two separate meanings that knowledge attributions carry, and then utilizes relevance concerns to explain why the second, pragmatic meaning comes to be focused on in the relevant examples.

Even where Rysiew's account is thus understood as simply positing two such separate meanings, I have serious worries about the security of the inference it ascribes to the hearers of the knowledge attributions, because in many such cases, including the high-standards Bank Case under discussion, the semantic content the account ascribes to the speaker's assertion *is* very relevant to the purposes at hand, even though the *issue* of whether the subject satisfies the pragmatic content the account ascribes is more conversationally relevant than is the issue of whether the subject satisfies the semantic content posited by the account. But we will take up such concerns below, when we are discussing the accurate understanding of Rysiew's proposal, and focus for now on whether the account understood as we are currently construing it satisfies our concerns about being properly based on general rules of conversation.

Note this crucial difference between Rysiew's WAM, so (mis-)understood, and the example of a plausible WAM we have looked at. The defender of the DKO account of possibility statements proposes an account of the semantic content of sentences of the form 'It is possible that p_{ind}', and then explains away the apparent counter-examples to the theory by means of that account of the sentences' meaning together with a very general rule of conversation ('Assert the Stronger'). That is, the defender of

---

[41] As Rysiew explains in both (2001) and (2005), though he himself adopts a relevant alternatives framework in presenting his invariantist account of the data often cited in support of contextualism, the basic thrust of that account does not really depend on that relevant alternatives framework, and can be adapted to a neutral, generic framework on which knowledge that p is simply understood as having a true belief that p and being in a good enough epistemic position with respect to p—where we remain neutral about what being in a good epistemic position consists in. Within such a generic framework, his account claims that the semantic content of 'S knows that p' is that S has a true belief that p and is in a good enough epistemic position to meet a particular standard, the standard for actually knowing, which, according to Rysiew does not vary from context to context. What the assertion pragmatically imparts within this generic framework is that S has a true belief that p and is in a good enough epistemic position to meet a contextually variable standard that depends on the particular conversational purposes that are in play as the assertion is made. Following Rysiew's own example, I will sometimes present his account in the relevant alternatives clothing in which Rysiew himself usually tends to put it, but my criticisms of the account will not depend on so understanding it, but will be applicable to its more generic version as well.

DKO shows how his theory of the semantic content of the relevant sentences predicts that assertions of the relevant sentences will, to use the terminology Rysiew works with, 'pragmatically impart' that the speaker doesn't know that the relevant p is true. By contrast, as we are currently understanding him, Rysiew helps himself to two different meanings that he simply posits that his target sentences bear—a 'semantic' and a 'pragmatic' meaning—and then, in the cases problematic to his semantics, uses a general rule ('Be relevant!') to explain how the pragmatic meaning comes to be the one that is focused on. So understood, Rysiew does not, as I urge, show how his proposed account of the sentences' semantics, together with general rules, predicts what is pragmatically imparted. Rather, he attempts to show how his proposed semantics, together with a general rule *and a second meaning for the relevant sentences that he simply helps himself to*, predicts that his proposed pragmatic meaning is what gets focused on by the listener and what will be primarily communicated by the speaker.

I am highly suspicious of accounts that help themselves to two such meanings in the way this one does, not only disliking the loss of economy in explanation, but also worrying that we will not be able to combat all manner of absurd theories about the truth-conditions of various sentences if defenders of those theories are able to posit separate pragmatic meanings that do the work of accounting for usage in the troublesome cases, allowing their account of the truth-conditions to sit safely off in the corner, out of the fray. Indeed, the semantic meaning posited by the theory, so understood, seems quite expendable to the explanation being offered—though of course there may be other reasons for positing such an invariant semantic content to the sentences. Since the account begins by helping itself to the pragmatic content anyway, it seems that it could just as well account for the uses in question by utilizing *only* that pragmatic meaning. Rather than saying there are two candidate meanings and utilizing relevance concerns to explain how the pragmatic meaning comes to be focused on, it could just simply propose that whenever assertions are made that carry such a pragmatic meaning (whether or not that's in addition to a separate semantic meaning), hearers focus on that pragmatic meaning.

Of course, in some sense, we all must hold that our assertions carry meanings beyond what goes into their truth-conditions. As we've seen, I hold that 'It's possible that p' conveys that the speaker doesn't know that p, though that thing conveyed is not part of the truth-conditions of the utterance. But I find such a claim unobjectionable when the second, pragmatic meaning is generated by the semantic meaning together with general principles. In such a case, it seems that the semantic meaning is responsible for handling the phenomena, working through the general principles. I am open to hearing evidence that certain assertions carry two meanings in the way posited by Rysiew's account, as we are currently misunderstanding it, but I think it is rational to prefer explanations that appeal to only one 'basic', we may say, meaning—the other, 'non-basic' propositions conveyed being generated by that one meaning and general principles.

But at (2005: 63–7; esp. 63–4), Rysiew explains that he did not intend his account to be understood as we have just been construing it. Rather, he intended the account to start out by positing only the invariant, semantic meaning it ascribes to the assertion: that the subject (has a true belief that p and) can rule out the relevant alternatives to p, where, again, which alternatives are relevant is not a context-sensitive matter. Rysiew intended to explain how, given that the assertion has been made with that content, together with general rules of conversation and other knowledge the hearers in the case can be presumed to have (as Rysiew put it: 'as well as parties' knowledge of the context, their general background knowledge, and so forth' (2005: 63)), the listener can be counted on to arrive at the conclusion that what the speaker intends to communicate is that the subject (has a true belief that p and) can rule out the *salient* alternatives to p, where, again, which alternatives are salient is a highly context-sensitive matter. Relevance concerns, then, are not intended as merely explaining how one of two already-posited meanings comes to be focused on, but, together with knowledge the subjects can be presumed to have, to actually generate the second, pragmatic meaning as something that the speaker can be inferred to be intending to communicate given that she asserted something with the invariant semantic content the account posits. Rysiew further elaborates that the semantic content his account assigns is intended to play 'an essential role in leading the hearer to the information [Rysiew] claim[s] is pragmatically conveyed, and not just in virtue of its (sometimes) irrelevance' (2005: 64). That the speaker asserted something with the invariant content Rysiew posits does not only initiate a search on the part of the hearer for what the speaker might be intending to communicate by his assertion, but is also supposed to *direct* the hearer's search. Because what the speaker has said is that the subject has a true belief and is in a good enough position to meet a particular standard for epistemic position (or, within the relevant alternatives framework: that he can rule out all the *relevant* alternatives), it's natural for the hearer to look for something closely related to that as what's really being communicated, and to therefore arrive at the conclusion that what the speaker intends to communicate is that the subject has a true belief and meets another (in this case, higher) epistemic standard (or that she can rule out all of the *salient* alternatives).

Understood in this way (the way Rysiew intended), then, the account certainly is of just the general type we are seeking. However, so understood, it seems the account cannot work. Though it seeks to provide the right type of account, it cannot adequately explain the phenomena.

Consider the case that Rysiew himself focuses on: my high-standards Bank Case. What the speaker says in that case is that he 'doesn't know' that the bank is open on Saturdays. As we are now (correctly) construing it, the only meaning Rysiew's account initially posits is its proposed semantic content for the assertion: that the subject doesn't satisfy the particular moderate standard that in fact is always the semantic standard for 'knowledge'; let's abbreviate that by saying that, on the account, the semantic content of the assertion is that the speaker doesn't know$_M$ that the bank is open on Saturdays. And on Rysiew's account, the subject does know$_M$ that the bank is open on Saturdays,

so the assertion is false.[42] However, on Rysiew's account, relevance concerns, together with knowledge the hearer can be presumed to possess, can be counted on to get the hearer from the falsehood that, according to the account, was semantically expressed by the assertion, to the truth that the subject does not know$_H$ that the bank is open on Saturdays—that he does not satisfy the higher epistemic standards that, according to Rysiew's account, the speaker pragmatically conveys.

But how could relevance concerns drive the hearer from the semantic content posited by the theory to the theory's proposed pragmatic destination? The falsehood that, according to the account, the speaker has asserted would be very relevant to the concerns of the participants in this conversation. We will grant that the *issue* or the question of whether or not the subject knows$_H$ is the most relevant issue to the purposes at hand (and is therefore more relevant than is the question of whether he knows$_M$): Whether or not he should go into the bank to brave the long lines (or at least verify by high standards that the bank is open on Saturdays) seems to stand or fall on the issue of whether he knows$_H$. But that the speaker doesn't even know$_M$ that the bank is open on Saturdays—which according to Rysiew's account is what the speaker has asserted—would of course settle (in the negative) the salient question or issue of whether he knows$_H$. So, in saying that he doesn't 'know' that the bank is open on Saturdays, on Rysiew's account, the speaker asserts what would be an extremely relevant thought. So, how can relevance concerns securely lead the listener to the conclusion that what the speaker intends to convey is not that extremely relevant thought he has asserted, but rather some other proposition?

Rysiew could appeal to—and at (2005: 65) he may be appealing to—the fact that the hearer in the high-standards Bank Case is aware of the facts in virtue of which the speaker knows$_M$ that the bank is open on Saturdays. And indeed, if it is common knowledge to the two of them that the speaker really does know$_M$ that the bank is open on Saturdays, then it is quite understandable that, when the speaker says that he *doesn't* know$_M$ that fact, the hearer would at least get quite confused, and might even start fishing around for something other than the semantic content of what the speaker said as an interpretation of the thought that the speaker was intending to convey—in which case that the speaker doesn't know$_H$ that the bank is open on Saturdays would seem to be a leading candidate. Thus, there is some hope for a pragmatic account in the vicinity of Rysiew's to work for the high-standards Bank Case as it is actually formulated, though it would not really be the *irrelevance* of what the speaker says that would tip the listener off that something else is being conveyed here, but rather the fact that it is common knowledge that what the speaker asserted is false.

However, it is incidental to the case that the hearer knows that the speaker knows$_M$ that the bank is open on Saturdays, and if we modify the case so that prior to the

---

[42] Rysiew chose to deny the truth of the denial of knowledge in the high-standards Bank Case. If he had instead decided that it was the positive claim to know in the low-standards Bank Case that was false, and had provided an account analogous to his actual account for how a relevant truth comes to be pragmatically conveyed by the false assertion there, my objections, suitably modified, would apply to the analogous account as well.

assertion the hearer has no idea of how well positioned the speaker is with respect to the bank being open on Saturdays—she is ignorant of whether he knows$_H$ and whether he knows$_M$—nothing like Rysiew's account seems to have any hope of working. The speaker's admission that he doesn't 'know' in this highly charged case, in which the possibility of the bank's having recently changed its hours has been explicitly raised and taken seriously, will seem as correct as it does in the unmodified version. But now it is extremely difficult to see what, on Rysiew's account, could be expected to drive the hearer from what the speaker has said to some other proposition as being what the speaker really intended to convey.

These *somewhat* analogous cases may make clearer the problem here.[43] Suppose first that we are discussing Sally's height, and, before you make the assertion we will focus on, I don't know anything substantial about Sally's height and I also don't know what, if anything, you know about it. And suppose that, in a way analogous to how the issue of whether the speaker knows$_H$ that the bank is open on Saturdays is the most conversationally relevant issue about his epistemic position with respect to that proposition in the Bank Case, the most conversationally relevant issue or question about Sally's height, given our current conversational interests, is whether she meets the rather high standard of being at least 6 feet tall: Given current interests, if I could learn the answer to just one question of the form 'Is Sally at least ___ tall?', I would ask whether she is at least 6 feet tall. (Perhaps we are discussing whether she is a good candidate to play center for a certain women's basketball team that plays at a level that makes 6 feet tall a good dividing line for our conversational purposes.) Now, first, suppose that you know that Sally is less than 5½ feet tall—and, of course, then, also that she is less than 6 feet tall. While it seems fine for you to say, 'Sally is less than 6 feet tall', it also seems perfectly fine (in fact, even better) for you to assert 'Sally is less than 5½ feet tall'. Saying that Sally is less than 5½ feet tall settles the highly relevant question of whether she is at least 6 feet tall, and also provides some further information about how far she falls short of that standard. If I realize that the most conversationally relevant issue about Sally's height is whether she's at least 6 feet tall, your asserting 'Sally is less than 5½ feet tall' certainly does not put me in any position to infer that you think that the most relevant issue about Sally's height is whether or not she's at least 5½ feet tall (I can perfectly well realize that you're aware that that's too low a standard to be the most relevant issue, given our current purposes), nor that what you have semantically expressed is conversationally irrelevant. Indeed, what you have semantically expressed is perfectly relevant: It settles (in the negative) the salient

---

[43] These cases are used just to clarify important points in my argument against Rysiew's account. Some of the disanalogies between these new cases and the relevant examples involving 'know(s)' are potentially important. I am not arguing here by analogy: I am not arguing that since the claim in the second 'tall' case we are about to encounter is wrong, then the denial of knowledge in the high-standards Bank Case would also be improper if that denial had the semantic content Rysiew assigns to it. Rather, I am only using these new cases to clarify the way in which the denial of knowledge in the high-standards Bank Case actually would be extremely relevant to the conversational purposes in play there, even if it were governed by only moderate standards.

issue of whether she's at least 6 feet tall, and gives further information about her height as well. There's no reason for me to think that the proposition you have asserted isn't just what you intended to convey.

Our second case will be just like the first (in particular, as before, we suppose that the most contextually salient issue concerning Sally's height is whether she is at least 6 feet tall), except that we will this time imagine that Sally is a bit taller than in the first case and that what you know about Sally's height is that Sally is between 5½ and 6 feet tall. Now, can you properly assert the falsehood that 'Sally is less than 5½ feet tall', relying on me to pass from the falsehood you have asserted to the true thought that Sally is under 6 feet tall as my conclusion as to what you were really trying to convey? Of course not. You cannot properly make that false assertion. From my point of view as the listener, this case is just like the previous one, and I will simply take you to be trying to convey just what you have said, since what you have asserted would settle (in the negative) the most salient issue concerning Sally's height, and would give further information as well. So, I will simply take you to be communicating that Sally is less than 5½ feet tall—which, in fact, is something you know to be false. No wonder, then, that in this case you cannot properly assert that 'Sally is less than 5½ feet tall'!

Likewise, given Rysiew's invariantist proposal regarding the content of the speaker's denial of knowledge in the high-standards Bank Case, it's hard to see how relevance concerns would drive the hearer to suppose that anything other than what was said is what the speaker intends to communicate (except in the special case where it's common knowledge that what the speaker said is false). That the speaker doesn't know$_M$ that the bank is open on Saturdays would settle (in the negative) the conversationally salient issue of whether he knows$_H$ that it's open, while going on to give more information as well, and is thus quite relevant indeed. At least where it isn't common knowledge between the speaker and the hearer that the speaker does know$_M$ that the bank is open on Saturdays, there doesn't seem to be any reason for the hearer not to suppose that the speaker simply meant to convey precisely what he said—that he doesn't know$_M$ that the bank is open on Saturdays. It's thus hard to see how Rysiew's account, properly understood, adequately explains what it sets out to explain.

Early Unger

Unger (1975) defended what we are here calling skeptical invariantism, holding that the truth-conditions for knowledge attributions are so demanding that our positive claims to the effect someone knows something to be the case are never true, and he sought to WAM the conversational data supportive of contextualism. While Unger did not appeal to a fully general conversational rule like 'Assert the Stronger' (as our defender of DKEW does), he did treat 'know(s)' as an instance of a *fairly* wide group of terms, which he called 'absolute' terms. A precise exposition of where and how Unger drew the line between absolute and other terms is beyond the scope of this short appendix. It will suffice for our purposes to note that it is indeed a fairly wide group of terms, which includes such words as 'flat', 'dry', 'straight', 'empty', and 'square'. (Unger's division of 'absolute' from other adjectives, and, based on that, his division of other absolute

terms—like the verb 'to know'—from non-absolute terms, can be quite important and helpful even if we end up rejecting Unger's proposed impossibly demanding semantics for absolute terms.) Thus, while Unger held the skeptical invariantist position that for any assertion of 'S knows that p' to be true, S must be in such a strong epistemic position with respect to p that it is impossible that anyone should be better positioned with respect to any proposition than S is positioned with respect to p, with the result that none or almost none of our attributions of knowledge is true, he did not thereby make 'know(s)' an 'isolated freak' of our language, as he put it (1975: 54). For he posited similarly stringent semantics for the many other absolute terms, according to which, in the case of 'flat', for instance, a surface had to be one flatter than which no surface can be conceived (to use some appropriate Anselmian wording to express Unger's position) for a sentence describing it as 'flat' to be true. In general, then, according to this view, positive assertions containing absolute terms (like 'S knows that p' or 'X is flat') have incredibly demanding truth-conditions, making all or almost all of our uses of such positive claims false. However, according to Unger's proposed account, when the person, object, or situation being described comes *close enough* for conversational intents and purposes to satisfying the demanding criteria for the true application of the relevant absolute term, it is warranted to falsely apply the absolute term—to say that the surface is flat, that the person knows, etc.

Unger's account, then, is not guilty of appealing to a special rule for the warranted assertability only of 'know(s)'. Consequently, Unger's defense of invariantism is on somewhat firmer ground than are the WAMs in virtue of which the positive case for contextualism is typically dismissed, at least with respect to our first (generality) desideratum. And Unger's proposed account has this going for it: If our language did somehow manage to develop a wide range of simple terms with impossibly demanding semantics like this, and yet we continued to use these terms to describe humble objects or situations which obviously didn't satisfy the semantic standards that govern the terms, some theory like Unger's would naturally suggest itself as an account of our use of these terms. But why suppose in the first place that the semantics of these terms actually is so demanding, despite our relaxed use of them?

Despite its not being an utter failure in terms of our generality considerations, several features of Unger's old invariantism work together to make it unattractive. First, note that while Unger's account doesn't posit a special rule, specific to 'know(s)', it also does not work by means of a fully general rule that applies to, and can be tested against, sentences of any content. By not utilizing a thoroughly general rule which has clearly correct applications, like 'Assert the Stronger', Unger loses a lot of leverage in advocating his view. His rule of assertability would be something like: 'When x is close enough, for present conversational intents and purposes, to satisfying the semantic requirements for "F", where "F" is an absolute term, then it is warranted, even if false, to describe x by means of F.' But a general contextualist account of the use of what Unger calls 'absolute' terms, which utilizes Unger's proposed context-variable warranted assertability condition as a truth-condition of the relevant assertions, and which thereby avoids the systematic falsehood that plagues Unger's general account,

will always be available, so it's difficult to see where the pressure to accept Unger's impossibly demanding invariantist account will come from. By contrast, a WAM that utilizes a rule that, like 'Assert the Stronger', is fully general and that has clearly correct applications, will have the decided advantage of being based on a rule that we are independently motivated to accept.

Next, recall that Unger's account falls on the bad side of our other two desiderata. And the move that allows Unger to do a bit better by our generality desideratum also exacerbates his failure with respect to the other features we look for, spreading the impossibly demanding semantics he posits and the resulting systematic falsehood in what we say to a large stretch of our language. Unger's invariantism is of the skeptical variety, and I suspect that most who reject the contextualist's varying standards don't imagine that the constant standards they will end up instead endorsing will be so demanding as to be unmeetable by mere mortals, and will be even more alarmed to learn that similar problems affect a wide swath of our ordinary talk.

Perhaps, however, a very general invariantism about absolute terms of at least roughly the kind Unger mounts could be used in defense of a moderate or non-skeptical invariantism, according to which the constant standards that govern the truth-conditions of sentences containing absolute terms will be meetably low. Such a moderate account will not be able to utilize Unger's account of assertability going by whether the sentences in question are *close enough* to satisfying the truth-conditions for those sentences to be appropriately asserted, for on such moderate accounts, the warranted assertability conditions for the sentences will often be more demanding than their truth-conditions. Defenders of such accounts could, however, posit both a moderate invariant standard to provide the truth-conditions for sentences containing absolute terms and add to that the claim that assertions of such sentences will be warranted if the object or situation being described has what is, for current conversational intents and purposes, enough of the property in question. Thus, 'X is flat' will be said to have moderate, invariant truth-conditions, but will also carry as its warranted assertability condition that X be flat enough, for current conversational intents and purposes; 'S knows that p' will likely be assigned moderate invariant truth-conditions, plus the warranted assertability condition that S be in an epistemic position that's good enough to satisfy current intents and purposes.

This moderate invariantist strategy, however, suffers from a disadvantage similar to one that plagues Rysiew's theory as I was initially inclined to (mis-)construe it. In fact, this just *is* Rysiew's account as I was initially inclined to construe it—except that this time the theory is being applied not just to 'know(s)', but to other absolute terms as well. This modification of Unger's account also does not derive the warranted assertability conditions it alleges from its account of the truth-conditions of the relevant sentences together with general rules, instead divorcing the semantics (truth-conditions) of the relevant sentences from their pragmatics: The moderate invariant semantic standard posited to provide a truth-condition for the sentences does no work in accounting for why and when those sentences are being appropriately and inappropriately asserted in various contexts. That work is all done by the quite independent warranted assertability

condition that the strategy simply posits. Having so divorced the semantics of sentences containing absolute terms from their pragmatics, the theory we are now considering offers us no good reason to accept the moderate invariantist truth-conditions for the relevant sentences that it posits. The work of accounting for ordinary usage is all done by the independent warranted assertability conditions the account posits. And the truth-conditions posited by the account don't do a very good job of yielding intuitively correct verdicts regarding the truth-values of knowledge-ascribing assertions made in various contexts. Its verdicts aren't as counter-intuitive as are those of skeptical invariantism, but they can't measure up to contextualism's performance here.

We have briefly looked at four invariantist theories here: (a) the moderate invariantist account that results from my earlier misreading of Rysiew, (b) Rysiew's own account, correctly understood, (c) Unger's skeptical invariantist position of (1975), and (d) a modification of Unger's position that takes a moderate, as opposed to a skeptical, approach to absolute terms generally. Again, all four completely fail to satisfy our second and third desiderata from section 11. Unger's skeptical invariantist theory does a bit better by our first test—that of generality. But its success there is quite limited, and comes at the substantial cost of charging a *whole lot* of perfectly appropriate, and seemingly true, claims made by competent speakers who are not basing their claims on any underlying factual errors with falsehood.

4

Single Scoreboard Semantics

1. Contextualism and Philosophical Debates over Skepticism

An important application of contextualism is to the problem of philosophical skepticism. The basic contextualist strategy is to explain the intuitive power of the skeptic's argument by claiming that the presentation of the argument triggers conversational mechanisms that have a tendency to raise the standards for knowledge to a level at which the skeptic's conclusions that we don't 'know' this or that are true. Thus, it is hoped, our ordinary claims to 'know' can be safeguarded from the apparently powerful attack of the skeptic, even while the persuasiveness of the skeptical argument is explained. For the fact that the skeptic can threaten to invoke extraordinarily high standards that we don't meet has no tendency to show that we don't satisfy the more relaxed standards that are in place in more ordinary conversations and debates.[1] Assessing the adequacy of such a response to skepticism will be one of the main objectives of Volume II.

Now, suppose a skeptic presents a skeptical argument to the conclusion that her conversational partner does not know something that she would ordinarily take herself to know, and that the argument being employed is one to which the contextualist applies the basic contextualist strategy described above. Suppose, however, that our skeptic meets with an 'Aw, come on!' response from her opponent, who insists, and continues to insist, that she does indeed know the thing in question. According to the contextualist, who 'wins'? Who is speaking the truth? The skeptic, her opponent, both, neither? What happens to the truth-conditions of knowledge-attributing and -denying claims in cases where the parties to the discussion seem to be pushing the 'conversational score' in different directions?

While many have presumed that contextualists would give certain answers to the above questions, the answers are in fact far from obvious. To answer

---

[1] For more on the basic contextualist strategy, see DeRose (1995: 4–7).

them one must take a position on what happens when different participants in a conversation are pushing the conversational score in different directions. In the work where I did most of my contextualist wrestling with skepticism—my (1995)—I remained neutral about such questions. Here, I'll present the answer I now favor and my reasons for favoring it, after laying out several other options, a couple of which are at least suggested by David Lewis in his very influential 'Scorekeeping in a Language Game' (1979), which, of course, inspired my use of 'scores' and 'scoreboards' in thinking of these matters. I'm very interested in what others think about this question: How other contextualists are inclined to answer the question, and what non-contextualists think is the contextualist's best answer. I hope that laying out the issue and some of the options will encourage some good thinking on this topic.

2. Contextualism and Disagreement

The objective described above argues for including this chapter in Volume II, where the contextualist treatment of skepticism will be developed and defended. This is instead here in Volume I because addressing the question of what happens to the truth-conditions of knowledge-attributing sentences in cases of at least apparently conflicting standards is also important to answering an important objection to contextualism, and therefore to making the case for contextualism.

When one speaker says that a certain subject 'knows' a certain proposition to be true, and another speaker says that that same subject does not 'know' that same proposition to be true, it will sometimes happen that the contextualist will hold that the two claims are compatible, because the claim of the speaker denying 'knowledge' is governed by higher standards than is the affirmation of 'knowledge', while the invariantist, of course, will hold that the claims are quite incompatible with one another. Where it intuitively seems clear that the claims of the two speakers are incompatible with one another, that contextualism rules otherwise constitutes an objection to contextualism. It's that objection that I address in this chapter. But are there such cases where contextualism issues a strongly counter-intuitive verdict of compatibility?

Well, when does contextualism actually rule that such a pair of claims are compatible? We've already investigated some cases in earlier chapters, especially Chapters 1 and 2. As we've seen, the main argument for contextualism is fueled by such pairs of claims. The contextualist not only accepts that the denial of 'knowledge' in HIGH and the affirmation of 'knowledge' in LOW are

compatible with one another, but also that they are both actually true, and uses the intuitive judgements that those assertions are true as premises in making the case for contextualism, where HIGH and LOW, recall, are the contextualist's cases, carefully constructed to render as plausible as possible the verdict that both assertions are true and that, therefore, the standards for knowledge are different in the two cases. While the contextualist's individual judgements about the truth of the two assertions involved are intuitively very plausible, it could turn out that it's also intuitively plausible that those claims are inconsistent with one another. If that's how things turned out, we would be faced with a conflict among three intuitions: that the attribution of 'knowledge' in LOW is true, that the denial of 'knowledge' in HIGH is true, and that the attribution of 'knowledge' in LOW is inconsistent with the denial of 'knowledge' in HIGH. However, as I will argue in the following chapter, we actually don't face much of a conflict here, because when you look at what HIGH and LOW are actually like, the claim that the assertions involved contradict each other is not very powerful at all—and cannot stand up against the contextualist's premises.

However, there are other cases where the intuition of incompatibility is quite powerful indeed, and such cases will be our focus in this chapter. Chiefly, as I've already indicated, we'll be looking at cases of dispute over whether or not something is 'known', whether such disputes involve philosophical skepticism or some more mundane issue over what is known. Recall that the two speakers involved in the contextualist's argument, S-HIGH and S-LOW, are not talking to one another, but are in completely different conversations. By contrast, in the situations we'll be mainly focusing on here, our two speakers are not only talking to one another, but are giving every indication that they mean to be, and take themselves to be, contradicting one another. In such cases, the sense that the two speakers are indeed contradicting one another can become very powerful indeed.

That could be a strong strike against contextualism if contextualism's verdict about such cases were that the denial of 'knowledge' in such a dispute is compatible with the affirmation of 'knowledge', and that both assertions are true. And many assume contextualists do render just such a counter-intuitive verdict about such cases of dispute. Here is an example of that assumption making its way into print; in a recent paper, Mark Richard writes:

Suppose a confrontation between a skeptic with high standards, and Moore, who has low standards. The skeptic says

(1) You don't know that you have hands.

Contextualism tells us that the content (and thus the extension) of 'knows' in the skeptic's context is determined by the standards that his context provides. Since he,

unlike Moore, has high standards, Moore and the claim that he has hands just don't make the cut. The skeptic's utterance of (1) is true: that is, Moore doesn't know that he has hands.

Of course, when Moore utters

(2) I know that I have hands

the standards in his context are the relevant standards, and so, given his low standards, he speaks truly. So Moore knows that he has hands after all. But how can that be? Didn't the skeptic just establish that Moore doesn't know that he has hands? Well, says the contextualist, what the skeptic said was true. But since 'know' is contextually sensitive, (2) doesn't say the same thing, when Moore uses it, as does

(3) You know that you have hands

when the skeptic uses it. So there's nothing contradictory about the skeptic's being able to use (1) truly while Moore can so use (2).

One feels that something is awry. One wants to say that when the skeptic and Moore argue with each other, they disagree about whether Moore knows that he has hands.   (2004: 215–16)

As I've already indicated, I agree that what Richard writes that 'one wants to say' here is extremely plausible. But does the contextualist really deny that? That is what we will be investigating here. It may well be that 'something is awry' here, all right—not just with the claim being attributed to the contextualist, but also with attributing such a claim to contextualism in the first place.

In addition to cases of speakers in such face-to-face confrontations, there are also cases of what I will call 'one-way disputes' where the intuition that the two speakers involved are making incompatible assertions can also get quite powerful. These will be addressed in section 12.

3. The Type of Debate Addressed Here

So, what does, or should, the contextualist say about the kind of impasse imagined by Richard between a philosophical skeptic and a Moorean opponent? Many seem to assume, like Richard, that the contextualist's answer will be that both the skeptic and her opponent are speaking the truth. The skeptic's claim that her opponent 'does not know' is true iff her opponent fails to meet the extraordinarily high standards that the skeptic's claim has a tendency to put into place. Since her opponent does not meet those standards, the skeptic's claim is true. Her opponent's insistence that she 'does know' is true iff she meets lower, ordinary standards for knowledge. Since she does meet

those lower standards, her claim is true, too. Both are speaking the truth, and they are failing to contradict one another. They are talking past one another.

There are certain ways the debate could go which I think would make that the correct answer to our question. For example, suppose the skeptic says, 'You don't know, and by this I mean…', and completes the sentence by explaining the very high standards her opponent would have to meet before the skeptic would count her as a knower. Her opponent replies, 'I do know, and by this *I* mean…', completing the sentence by specifying the moderate epistemic standards that she is claiming to meet. Both speakers having specified what they mean, and having specified very different meanings, they stop thinking of themselves as contradicting one another. Even if I were an invariantist, under these circumstances I would think that both the skeptic and her opponent are speaking the truth. As an invariantist, I'd think that the truth-conditions of standard uses of 'know(s)' do not vary with context. But when a speaker resorts to an explicit use of a 'and by this I mean…'-like construction, I'm inclined to think that the speaker's utterances come to mean what the speaker says they mean. Speakers are free to stipulate what they will mean by a new term they are introducing, and they're free to stipulate a special meaning they are giving to an existing term. Thus, if I were an invariantist, I'd think that in the imagined situation, at least one of our two speakers is using 'know(s)' in a special, non-standard way, and that both the skeptic and her opponent are speaking the truth. (But I would also think that such a conversation shows nothing about standard uses of 'know(s)'.)

Different situations that do not include anything that looks like a stipulation of (perhaps non-standard) meaning also tempt me, as a contextualist, toward the same verdict. Suppose, for instance, as often happens in such debates, that one or the other speaker, or some third party, advances some contextualist, or proto-contextualist, analysis of the debate that is transpiring, phrased in terms of what each party to the debate 'means' by 'know(s)', saying something along the lines of, 'Well, by "knows", you seem to mean, … while I seem to mean…'. Suppose those involved in the debate all accept this analysis, and our speakers, now believing that they mean different things by 'know(s)', consequently stop thinking that they are contradicting one another. I won't venture a guess as to what invariantists will say about a case like this, where the participants to a discussion have explicitly accepted such an analysis of their own debate, but have not said anything that looks like an explicit stipulation of meaning. But as a contextualist, as I've admitted, I'm at least tempted to conclude that both the skeptic and her opponent are speaking the truth in such cases. But that is not the kind of case I'll be addressing in this chapter.

It often happens in *disputes* between skeptics and their opponents that nobody even offers such an analysis of what's transpiring, much less is such an analysis accepted; the skeptic and her opponent *take themselves* to be contradicting one another; each *intends* to be contradicting what the other is saying; and, beyond what's going on privately in their own minds, each is publicly *indicating* that they are (or at least mean to be) contradicting the other, by saying such things as: 'No, you're wrong. I *do* know.' And similar situations arise in cases of non-philosophical 'skeptics' and their opponents. It's such cases as these that I'm addressing in this chapter. And about this kind of case, I am *not* inclined to think that both our speakers are speaking the truth, but failing to contradict one another.

The type of case in question is one in which there is a conversation crucially involving a context-sensitive term (one that can express different specific contents, given context) and in which, as we can put it, the *personally indicated content* of one speaker—the content that speaker's conversational maneuvers have a tendency to put into place for that term—diverges from the personally indicated content of the other speaker, but in which the speakers still indicate that they are contradicting one another.[2] This is a type of situation that can arise with other context-sensitive terms, too, and so, though I'm here mainly addressing what to say about such situations where 'know(s)' is involved, the question of what a contextualist should say about such a case is a quite general one.

I should note that I won't be getting into any specifics here about the precise mechanisms by which speakers can change the score or resist such changes. With regard to 'know(s)', different contextualists have proposed different accounts of the rules by which the standards for knowledge can be raised, and, though we haven't been specific about the precise means, we all seem to suppose that there are ways that standards can be lowered and ways that raises in epistemic standards can be resisted. We will just suppose our

---

[2] It is tempting—indeed, almost irresistible—to describe the situation as one in which the skeptic is *trying* to raise the standards for knowledge, while her opponent is *trying* to keep the standards low. But since what I am here calling the standards for knowledge are the standards a subject must meet for her to be *truthfully* credited with 'knowledge,' our debaters need not even be thinking of the standards, so construed, as being something that can be changed. They can be operating under the invariantist assumption that the truth-conditions for knowledge sentences remain fixed. Thus, they shouldn't be said to be trying to have an effect on the standards, so construed, which is why I instead described them as executing certain conversational maneuvers which have tendencies to put (or keep) in place certain standards. Typically, though, our debaters will be trying to change or affect *other* standards: the standards for when a subject will be accepted as 'knowing' in their conversation, for instance. But they may conceive of themselves then as trying to get those latter standards in line with the 'true' standards—those that articulate the conditions under which a subject can be *truthfully* said to 'know'—as trying to get their opponent to admit the truth of the matter under discussion.

skeptic has executed a maneuver (whatever that maneuver is and however exactly it works) that has a tendency to raise the epistemic standards, and that her opponent has responded by executing a maneuver that has at least some tendency to keep lower, ordinary standards in place. We are then concerned with the question of what happens to the conversational score in such a situation. For simplicity, I will (except for some brief speculation in n. 10) concentrate only on cases in which there are just two speakers involved in a conversation. Also, I will ignore complications caused by vagueness in the personally indicated content of certain speakers, and suppose that each speaker is personally indicating a precise set of standards.

4. Multiple, Personal Scoreboards

Many seem to think or assume that, even with respect to the cases just described, the contextualist's answer, or the contextualist's best answer, to our question will still be that both of our debaters are speaking the truth, and they are failing to contradict one another. The idea here is that the truth-conditions of each speaker's spoken claims will directly match the personally indicated content of that speaker. Thus, since the skeptic is executing conversational maneuvers that tend to put into place the high, skeptical standards for knowledge, her claims that her opponent 'doesn't know' are true iff her opponent fails to meet those extraordinarily high standards. And since her opponent's conversational maneuvers tend to put (or keep) in place lower, more ordinary standards, *her* claims that she does indeed 'know' are true iff she meets those lower standards. The picture seems to be that each speaker, in addition to having her own personally indicated content (having certain standards that her conversational maneuvers have at least some tendency to put in place) also has her own personal scoreboard, by which I mean that the truth-conditions of each speaker's use of 'know(s)' are particular to that speaker, and presumably match that speaker's personally indicated content.[3] Of course, it helps various communicative purposes if speakers engaged in a conversation adjust to one another's usage, and come to have matching scores on their scoreboards. And those who think in terms of multiple, personal scoreboards will probably think that's what usually happens. But in the cases we're considering here,

---

[3] Here my use of the metaphor 'scoreboard' differs from Lewis's. For Lewis, there is a scoreboard 'in the head' of each of the participants, and what the score is can be a function in part of what all these different scoreboards say the score is. As I use 'scoreboard' here, it by definition gives the right score.

that doesn't happen; our speakers fail to adjust to one another. And if you're a contextualist who is thinking in terms of each speaker having their own personal scoreboard, you will conclude that both our skeptic and her opponent are speaking the truth, and are failing to contradict one another.

This position, of course, does have its appeal. There is something to be said for thinking the truth-conditions of a speaker's knowledge claims match the epistemic standards that she herself is indicating by her conversational maneuvers.

5. Single Scoreboard Semantics

But there is a cost, too. Each speaker, in addition to indicating certain epistemic standards, also indicates that they are contradicting the other speaker. And the multiple scoreboards position has the unfortunate result that our two speakers are not, as they surely seem to be, and as they take themselves to be, contradicting one another.

One thought that can lead one to the multiple scoreboards view is that we should respect what each speaker is indicating in assigning content to the claims of that speaker: speakers should be in control of their own meaning. But, as I've just noted, in the cases in question, each speaker is indicating two different things: that the standards be such-and-such, and that their claims contradict those of the other speaker. But where the indicated standards of the two speakers diverge, we cannot consistently respect *all* of the indications being given by our speakers. Why should it be the clear indications being given by each speaker that her claims be understood to contradict those of the other that give way here?[4]

My own thought about the cases under consideration has always assumed what I call 'single scoreboard semantics', due to the influence of David Lewis's 'Scorekeeping in a Language Game' (1979) which seems to me to promote such an account.[5] On this view, there is a single scoreboard in a given conversation; the truth-conditional content of both our speakers' uses of 'know(s)' is given by the score registered on this single scoreboard. The score of course can

---

[4] The need to stress the point of this paragraph was pointed out to me by Robert Stalnaker.

[5] Due to his different use of 'scoreboard,' described above in n. 3, Lewis would not describe the position he tends toward as one in which there is a single 'scoreboard': He uses the term so that each speaker does have their own scoreboard in their head. However, Lewis does seem to assume in his writing that there is a single (though changeable) conversational score in a given conversation, rather than writing as if there is a different score for each speaker, and thus seems to be working with the picture that in *my* use of the term has a single 'scoreboard'.

change as the conversation progresses, and the score it registers is responsive to the maneuvers made by all the various speakers, but there is at any given time a single score that governs the truth-conditions of all the speakers' uses of the relevant term. Thus, both our skeptic's and her opponent's use of 'know(s)' are governed in their truth-conditions by what that single scoreboard registers when their claims are made. This picture, of course, promotes the thought that our skeptic and her opponent really are contradicting one another: If there is a single scoreboard, then our skeptic is denying precisely what her opponent is affirming as they debate back and forth.

Of course, if we imagine the score moving sharply up and down throughout the conversation—moving up (the standards for knowledge going up) every time the skeptic speaks, and dipping down suddenly whenever her opponent makes a claim—the result would be equivalent to supposing that each speaker had their own personal scoreboard. So that's not the idea. There are tricky questions we'll take a brief look at toward the end of section 8 about the exact timing of certain changes of score which can make it difficult for the contextualist who embraces single scoreboard semantics to know what to say about the truth-conditions of some of the opening claims made in our debate, and to know what to say about the relation between the contents of these early claims and later assertions, intended to be in opposition to the earlier claims. In section 13, I will make a suggestion about how to handle these issues. But, in any case, once the skeptic has made her standards-raising maneuvers a time or two, and her opponent has responded with her stubborn, resistant maneuvers, and they continue to debate—'You don't know', 'I do know'—the conversational score has presumably reached whatever state of equilibrium it reaches in such situations, and, at least as far as the truth-conditions of their knowledge claims go, on the single scoreboard picture, the skeptic is denying precisely what her opponent is claiming. So they are contradicting one another.

But which is speaking the truth? What *is* the conversational score in such a situation? Let's canvass some possible accounts of what happens to the conversational score—what happens to the truth-conditions of claims about what is and is not 'known'—in the situations we're considering.

6. Higher Standards Prevail, So the Skeptic Wins

In my experience, the view that is most often ascribed to contextual-ists—because it is thought this is what we do think, and/or because it is

thought that this would be the best view available to us—is the 'multiple scoreboards' view that both the skeptic and her opponent are speaking the truth. As I've already noted, I do not in fact accept this view. I've also asked Stewart Cohen, who has informed me that he doesn't accept that answer either. Nor does either of us accept the view that is the next most commonly ascribed to the contextualist: that, in the situations under consideration, the skeptic's extraordinarily high standards prevail, so that the skeptic is speaking the truth when she charges, 'You don't know,' and her opponent is saying something false when she insists, 'I do know.'

Some of what David Lewis wrote suggests this very skeptic-friendly view. Background: Lewis worked within the relevant alternatives picture of knowledge, according to which an attribution of knowledge of a proposition p to a subject S requires for its truth that S can rule out or eliminate all the relevant alternatives to p. The skeptic changes the score by enlarging the range of alternatives to p that are relevant so that they include various hard-to-eliminate alternatives that are usually irrelevant. In his (1979), Lewis writes:

The commonsensical epistemologist says: 'I *know* the cat is in the carton—there he is before my eyes—I just *can't* be wrong about that!' The sceptic replies: 'You might be the victim of a deceiving demon.' Thereby he brings into consideration possibilities hitherto ignored, else what he says would be false. The boundary shifts outward so that what he says is true. Once the boundary is shifted, the commonsensical epistemologist must concede defeat. And yet he was not in any way wrong when he laid claim to infallible knowledge. What he said was true with respect to the score as it then was.   (1979: 355)

It isn't completely clear what Lewis meant by saying that the commonsensical epistemologist 'must concede defeat'. What if, as in the situations we are here considering, the epistemologist doesn't do what he 'must' do, and stubbornly repeats his claim to know? It's tempting to (and I'm inclined to) read Lewis here as holding that the commonsensical epistemologist would then be saying something false, because, in this context, I am tentatively inclined to understand 'the commonsensical epistemologist must concede defeat' to mean that, if he says either that he 'knows' or that he doesn't 'know', he must choose (admit) the latter if he is to speak the truth.[6]

---

[6] One could quite plausibly understand Lewis differently here. For instance, if the force of Lewis's 'must' is understood differently, this passage could be read as being compatible with the multiple scoreboards view that both the skeptic and her opponent are speaking the truth. Perhaps if Lewis's 'commonsensical epistemologist' were to insist that he does 'know,' his claim would have the low,

In his later paper 'Elusive Knowledge' (1996), Lewis seems to accept an analysis that also yields that result. In that paper, the irrelevant alternatives are the ones that are 'properly ignored'; the relevant alternatives that must be 'eliminated' by one's evidence if one is to count as a knower, then, are those that are either not ignored, or are ignored but improperly so. The skeptic, by calling attention to various skeptical hypotheses, creates a situation in which those uneliminated possibilities are not ignored at all, and so are not properly ignored.[7] If the skeptic's opponent were to claim to know, then, it seems, he would be saying something false.

Perhaps Lewis's apparent backing of the verdict that the skeptic wins is at least largely responsible for its popularity. Another possible contributing cause is that other contextualists, as well as Lewis, have concentrated our efforts on discerning the mechanisms by which epistemic standards are raised, and have not worried that much (at least in print) about how standards can be lowered, or about how threatened raises in standards can be resisted.

7. Does It Matter if the Skeptic 'Wins'?

One common complaint against the contextualist approach to skepticism is that it is too friendly to the skeptic. To some extent, that complaint is unavoidable. For just about any contextualist will hold that we don't know according to the epistemic standards that the skeptic's maneuvers have at least some tendency to install, and that they actually succeed in putting in place when their moves are not resisted, and that by itself will be too soft on skepticism for the tastes of some. But the views we have considered so far make the situation worse for the skeptiphobic. For on the multiple scoreboards view, the skeptic truthfully states her conclusion even when she is resisted—though her opponent is also

---

ordinary truth-conditions, and would therefore be true, according to Lewis. The sense in which he 'must' concede defeat, on this reading, is that if he is to speak *properly*, he must not make such an ordinary, low claim in response to the skeptic. Still, on this reading, a defiant response would be, though improper, true.

[7] There are ways to use 'ignore' so that one can be said to be 'ignoring' something that in some sense one is attending to: In a philosophical discussion, I am paying very careful attention to a certain objection that has just been raised, planning to craft a response later, but I am properly described as 'ignoring' it as I proceed, in that I don't alter what I am now saying as a result of it. Similarly, if a skeptic brings to our attention a certain possibility of error, it seems to me that in a very good and relevant sense, we can choose to 'ignore' it, even as we are thinking about it, by, for instance, treating it as something not worth worrying about. But in his (1996), Lewis seems to use 'ignore' in such a way that if you are at all attending to a possibility, you are thereby not ignoring it.

held to be speaking truths when she insists that she does 'know'. And on the only single scoreboard view that we have looked at so far, things are even worse: Not only is the skeptic speaking the truth, but her opponent is speaking falsely.

That the two at least fairly skeptic-friendly views we have looked at to this point are so often assumed to be what the contextualist will say has probably made contextualism unattractive to many, and it might help make the contextualist case, at least to many, to point out that there are other, less skeptic-friendly, ways to take the contextualist approach.

But before considering these other options, including the one that I favor, I'd like to register my view that, though I think the skeptic does not 'win' in the situations we are considering—I don't think the skeptic is speaking the truth when she says we do not 'know' in situations in which she is resisted by a stubborn opponent—it would not matter that much if she did win. Once we see that we don't know according to the standards the skeptic's maneuver has at least some tendency to put in place, but that we do know according to the ordinary standards that govern most of our thinking and speaking about what is and is not 'known', so that the skeptic's 'success' has no tendency to show that we're usually deeply mistaken in our ordinary thought about what we do or do not 'know', we have seen most of what's important in the contextualist approach. We can then go on to discuss the somewhat important matter of how serious or worrisome a result it is that we don't meet the skeptic's extraordinarily high epistemic standards, once we're freed from the concern that that shows our ordinary thought about 'knowledge' to be mistaken. (My own view here is that this is not very worrisome at all, but I will leave the task of explaining why to Volume II.) But once all the above is seen, I myself don't find it all that important who 'wins'—who ends up speaking the truth—in various debates between skeptics and their opponents in which the parties do not see the contextualist resolution to their debate, and fail to adjust to one another's standards.

But those who see things differently, and are especially bothered by the thought the skeptic 'wins' debates that go the way we're considering here, may be especially interested in some of the less skeptic-friendly options we'll look at below. (For those who, like me, don't much care whether the skeptic 'wins' in the sense in question, this can still be an interesting question about how to understand conversations involving divergence in personally indicated content in context-sensitive terms.) We will start with a view that is very *un*friendly to the skeptic.

8. Veto Power

As I've noted, Lewis's contextualist writings tend toward the view that, in the situations we are considering, the skeptic wins. But in a section of 'Scorekeeping' different from the one we've looked at already, Lewis at least suggests a different view—not about knowledge, but about 'standards of precision':

> Taking standards of precision as a component of conversational score, we once more find a rule of accommodation at work. One way to change the standards is to say something that would be unacceptable if the standards remained unchanged. If you say 'Italy is boot-shaped' *and get away with it*, low standards are required and the standards fall if need be; thereafter, 'France is hexagonal' is true enough.   (1979: 352, emphasis added)

Here Lewis is explaining 'rules of accommodation': rules by which the meaning of context-sensitive terms tends to shift so as to make what is said true. Important for our current purposes is the phrase I've emphasized. What if our speaker does *not* get away with lowering of the standards of precision? Then, it at least seems to be suggested, the standards do not fall. Our speaker has said something that—in this case, by a 'rule of accommodation'—has a tendency to lower the standards of precision so that Italy counts as 'boot-shaped'. The picture at least suggested, though, is that those he is talking with have 'veto power' over this changing of the conversational score: If they don't let him get away with changing the score, then he does not succeed in changing the score.[8]

A similar position is tempting about our debate between the skeptic and her opponent. According to the contextualist, the skeptic has executed maneuvers that have a tendency to raise the standards for knowledge. But in the situation we are considering, her opponent does not accede to this raising of epistemic standards; the skeptic is not allowed to 'get away with' changing the score. On the view under consideration, then, the standards are not raised. Here, we get a view on which the skeptic's opponent ends

---

[8] Note, however, that Lewis's example involves the vetoing of a 'proposed' *lowering* of the standards of precision. Perhaps the way to put Lewis's suggestion of veto power together with his general friendliness to skeptical standards is to understand him as holding that conversational partners have veto power only over each other's attempts to lower standards. (Or perhaps it's just *easier* to veto the lowering of standards.) Whatever is the best reading of Lewis's own position here, I am in the text construing the 'veto power' option as the position that conversational partners have veto power over all potential changes to the conversational score—over both the raising and the lowering of epistemic standards, in the case of knowledge attributions. When this veto power is exercised, the conversational score stays where it was.

up speaking the truth, and the skeptic's claims are false. This depends on the assumption that the epistemic standards start out meetably low, perhaps because they have such a 'default' value, but I suppose that is a plausible enough supposition.

There is a tricky issue about timing here. Suppose our skeptic has just made her standards-raising maneuver, but her opponent has not yet responded. Her opponent is quickly deciding between two responses, one of which would constitute letting the skeptic get away with raising the standards, the other of which would constitute not letting her get away with it. But right now, in the brief moment before the skeptic's opponent speaks, where are the standards? Was the skeptic's knowledge-denying utterance true? Can the truth-value of that claim depend on something—how her opponent will respond—that has not yet occurred? Somewhat similar questions can arise for other views we will be looking at. I will make a suggestion as to how to answer these questions in section 13. In the meantime, we will assume that at least after our skeptic and her opponent have each had a couple of turns, the conversational score has reached whatever state of equilibrium it reaches in such cases of non-cooperation. According to the view currently under discussion, since the skeptic has failed to get away with raising the epistemic standards, and has therefore failed to raise them, it is her opponent who is speaking the truth when she continues to insist that she does 'know', while the skeptic is speaking a falsehood when she continues to claim that her opponent does not 'know'.

9. Reasonableness Views: Pure Reasonableness and 'Binding Arbitration'

To some, the matter of what would be a reasonable score, given the situation of a conversation, is important in determining what the score actually is in cases like those we are considering. The simplest way for this to work is what we can call the 'pure reasonableness' view, on which the conversational score—the relevant aspect of the truth-conditions of the various speakers' claims—is whatever would be the most reasonable score for the speakers to use, no matter what the personally indicated content of the various speakers may actually be.

Another way reasonableness can figure in is on what we may call the 'binding arbitration' model. In major league baseball, at least as I understand the process, when a player and his team go to binding arbitration to decide

the player's salary, each makes a bid or proposal on what the player's salary should be, the player, of course, suggesting a higher figure than the team. Evaluators then decide which of the two bids is the more reasonable, given the abilities of the player and the salaries of other players at the same position, and other factors, and that becomes the player's salary. On the 'binding arbitration' model of conversational scorekeeping, when the personally indicated content of one speaker diverges from that of the other that she is talking with, then the score—which gives the truth-conditions of both speakers' claims—matches the personally indicated content of the speaker who is indicating the more reasonable content of the two.

The binding arbitration view has some attractions. On it, as on other single scoreboard views, our skeptic and her opponent really are, as they seem to be, contradicting one another. And the personally indicated content of one speaker—the one that is indicating the more reasonable standards—matches the truth-conditions of her own claims. Of course, the truth-conditions of the other speaker's claims don't match up with her own personally indicated content, but that can seem like just punishment for her putting in the more unreasonable 'bid'.

On both of the reasonableness views we've considered, the matter of which of our speakers is speaking the truth depends on facts about what are the reasonable standards for them to use in their situation. And, of course, this opens up a whole host of questions that I won't even begin to address about what makes standards the reasonable ones to use. But anti-skeptics who think that the skeptic's standards for knowledge are entirely unreasonable and also think that we do meet any reasonable standards for knowledge will see in these views a way to deprive the skeptic of victory.

For my part, even if I could see past the problems involved in deciding what are reasonable standards, I wouldn't be attracted by reasonableness views of the type we have been considering. It's good for speakers to use reasonable standards (or, more generally, reasonable scores), of course. But if they opt for unreasonable standards, I'm inclined to think the truth-conditions of their claims then reflect those unreasonable standards that they are indicating. In uncooperative conversations where the different speakers are personally indicating different standards, if the indicated standards of one of the speakers are more reasonable than the other's, then the more reasonable speaker is in some way more conversationally praiseworthy than the other, I suppose. However, I don't see that this is so in any sense that would mean that the truth-conditions of both speakers' claims should be thought to have the truth-conditions that match the more reasonable speaker's indicated standards.

10. The Exploding Scoreboard

We have not yet arrived at the view I prefer. To prepare for it, we'll first consider another view, that, having grown up as a sports-crazy kid on the South Side of Chicago during the days when Bill Veeck owned the White Sox and Comiskey Park, I can't resist calling the 'Exploding Scoreboard' view.[9] On this view, when the personally indicated content of two speakers in a single conversation diverges—or at least when they diverge by as much as happens in our debate with the skeptic and her opponent—the scoreboard explodes: There is no correct score, and claims involving the relevant term are neither true nor false.

We are, after all, considering a conversation that is quite defective. Speakers engaged in conversation should adjust to one another's score so they are meaning the same thing by the key terms in question. When this doesn't happen, and especially when the divergence in personally indicated content is great, as in the situations we are here considering, perhaps semantic hell breaks loose.

On this view, neither the skeptic nor her opponent is speaking truths in the situations we are considering—nor is either speaking falsehoods. I'll let you decide for yourself whether and to what extent the skeptic 'wins' on this view—and whether that's a problem. She doesn't truthfully claim that we don't 'know', but her destructive purposes are served by her creating a context in which her opponent cannot truthfully claim that she does 'know'.

The Exploding Scoreboard view, initially at least, strikes me as very plausible when it's applied to cases of great divergence. It's far less plausible when applied to cases of slight divergence.

Consider an example involving the matter of how far out 'here' reaches in a given case. You and a colleague are in a giant room—say the grand ballroom of a hotel at the APA 'smoker'—and are discussing whether or not Frank is 'here'. By 'here', neither of you means to be designating an area so small that Frank would count as 'here' only if he were sitting right at the same table that you and your colleague are sitting at. (It's obvious that he's not right at your

table, so there would be no point in wondering, as you are, whether Frank is 'here' if that's what 'here' designated.) But neither do either of you mean to designate an area so large that the whole ballroom counts as being 'here'. (You both assume that Frank is somewhere in the ballroom, and so wouldn't be questioning whether Frank is 'here' if 'here' included the whole room.) You both mean to designate—and make conversational maneuvers indicative of this—something in between those two extremes. But the area that you persistently indicate is slightly larger than the area your colleague indicates, and you fail to adjust to one another's usage. (Don't ask me how this—one of you indicating just a slightly larger area than the other—could happen. Just suppose it does.) While this conversation is slightly defective (the two of you should somehow adjust to another's usage so that you mean the same area by 'here'), it's not all that bad, and if Frank is in fact in the far corner of the ballroom, clearly outside of what either of you have been indicating as the range of 'here', then it seems very plausible that in your conversation with your colleague, 'Frank is here' is false, and 'Frank is not here' is true. Thus, it is very implausible to suppose that, because of the slight divergence in personally indicated content in your conversation, the 'here' scoreboard has simply exploded.

11. The 'Gap' View

The single scoreboard view that cases of small divergence strongly suggest to me is the 'Gap' view, according to which, as applied to your conversation about whether Frank is 'here':

- claims that 'Frank is here' are true (and claims that 'Frank is not here' are false) iff Frank is in the region that counts as 'here' according to both speakers' personally indicated content;
- claims that 'Frank is not here' are true (and claims that 'Frank is here' are false) iff Frank is in the region that does not count as 'here' according to either speaker's personally indicated content; and
- claims that 'Frank is here' and that 'Frank is not here' are neither true nor false if Frank is in the region that counts as 'here' according to one speaker's personally indicated content, but does not count according to the other's.

Applied to cases involving 'know(s)', then, in cases of small divergence in personally indicated epistemic standards, 'S knows that p' is true (and 'S doesn't

know that p' is false) where S meets the personally indicated standards of both speakers; 'S doesn't know that p' is true (and 'S knows that p' is false) where S fails to meet either set of standards; and where S meets one set of standards but fails to meet the other, both 'S knows that p' and 'S doesn't know that p' go truth-value-less.

There's a lot to be said for giving a uniform treatment to cases of small and large divergence—especially since it will be difficult to draw a line between the two. So, having accepted the Gap view for cases of slight divergence, I'm inclined to accept it as well for cases of great divergence. And why not? Even where the divergence is great, isn't it plausible to suppose that 'S knows that p' is false if S fails to meet the personally indicated standards of either speaker? Just so, even in our debate between the skeptic and her opponent, it seems plausible that either speaker would be saying a truth if they were to claim that 'Frank does not know that the Cubs won the 1908 World Series' if in fact Frank fails to meet even the standards for knowledge of the skeptic's opponent.

Indeed, the Gap view is the view I'm most inclined to accept about the situations of dispute we have been addressing here since section 3. Its main attraction to me is its impressive ability to simultaneously respect both the sense that our two speakers are contradicting one another and the feeling that the truth-conditions of each speaker's assertions should match that speaker's personally indicated content.

The view does an excellent job of respecting the sense of contradiction—as well as can be hoped for in cases involving gappy claims. Suppose the person you're talking with makes a statement with 'gappy' truth-conditions. Take, for example, a standard case of vagueness. She says 'Frank is tall', where her claim is true if Frank's height is in a certain range, is false where Frank's height is in a certain lower range, and, we'll suppose, is neither true nor false if Frank's height is in an intermediate range. You reply, testily, 'Frank is not tall!' What content can we assign your claim that would do the best job possible of making your claim contradict your friend's claim? Seems to me, that's done by supposing that your claim is true in the same range of cases where your friend's claim is false, that your claim is false in the same range of cases where your friend's claim is true, and that your claim fails to have a truth-value in the same intermediate range where your friend's claim has no truth-value. Of course, we are then saying that in the intermediate range of cases, both you and your friend's claim match in that they both fail to have a truth-value. But this seems to me the position that does the best job of making you contradict your friend. You seem to be denying the very same gappy thing that your friend is affirming. In any case, where either of your claims have truth-values,

they both have truth-values and the truth-value of each is the opposite of that of the other.

On the Gap view, we get the same relation between our two speakers' claims in cases of contextually sensitive terms where there is a divergence in personally indicated content, where one speaker is asserting 'p' and the other 'not-p'. In the debate we have been imagining, our skeptic is denying the same gappy thing that her opponent is affirming.

What of the relation between the truth-conditions of each speaker's utterance and the personally indicated content of that speaker? Here, the Gap view does not deliver the closest match possible. Suppose you assert, 'Frank knows that the Cubs won the 1908 World Series', where the epistemic standards you are personally indicating are different from—let's say lower than—those that your conversational partner is personally indicating, in a case where the two of you fail to adjust to one another's usage. We *could* suppose here that the truth-conditions of your claim exactly match your own personally indicated content—that your claim is true iff Frank meets your personally indicated standards and is otherwise false. Indeed, the multiple scoreboards view delivers this exact match. The Gap view does not. On the Gap view, in a certain range of cases—where Frank meets your, but not your partner's, personally indicated standards—your claim is neither true nor false, though it would be true if the truth-conditions of your claim simply followed your own personally indicated standards. But this difference is unavoidable on any view that, like the Gap view, has you asserting exactly what your partner is denying. Where there is a divergence in personally indicated standards, no consistent view can make the truth-conditions of both speakers' claims match their own personally indicated standards while still delivering the result that one is denying precisely what the other is affirming.

However, while the Gap view doesn't do the best job possible of matching the truth-conditions of each speaker's claims with that speaker's own personally indicated content, it does provide for what seems to me a nice intimate relation between the two: In no possible case is your claim true where it would be false if it bore the content you personally indicate, and in no possible case is your claim false where it would be true if it bore the content you personally indicate. Alternative formulation: Wherever it is both the case that your claim has a truth-value, and that it would have had a truth-value if it bore the content you personally indicate, the truth-value it has is the same as the truth-value it would have had if it bore the content you personally indicate. Though, on the Gap view, your claim can sometimes fail to have a truth-value when it would have had a truth-value if its truth-conditions just were those that you personally indicate, it will never happen that your claim takes the opposite

truth-value from what it would have borne if it had as its truth-conditions what you personally indicate.

We might go so far as to call this relation that, on the Gap view, the truth-conditional content of your claim bears to your personally indicated content that of 'weak equivalence'. It's the same fairly intimate relation that I think 'p' bears to 'It is true that p'. Where 'p' is true, so is 'It is true that p'; where 'p' is false, so is 'It is true that p'; but where 'p' is neither true nor false, then it seems to me, as it seems to many others, 'It is true that p' is false:

p	It is true that p
T	T
F	F
N	F

In no case is one of 'p' and 'It is true that p' true while the other is false, or, alternatively, wherever they both have a truth-value, it is the same truth-value. On the Gap view, the truth-conditional content of your claims bears that same relation to your own personally indicated content.

The Gap view achieves this, to my thinking, very impressive result of both going a long way toward respecting both the sense of contraction and also securing a nice relation between the truth-conditions of each speaker's claim and that speaker's personally indicated content only by assigning a gap in the truth-conditions of the relevant claims, and in cases of great divergence in personally indicated content, this gap in the truth-conditions—the range of cases where the claims are neither true nor false—gets very large. So the view is likely attractive only to those who are not great fans of bivalence. But I, for one, don't mind gaps. And in very defective conversations, where there is great divergence in personally indicated content, I don't find it implausible to suppose that those gaps get very large. In fact, that seems to me the right verdict to reach about such conversations. After all, I found the Exploding Scoreboard view, on which the gap becomes all-engulfing, initially plausible as applied to cases of great divergence (though in the end I prefer the Gap view).

So, somewhat tentatively, I accept the Gap view. On it, as on the Exploding Scoreboard view, neither the skeptic nor her opponent is speaking the truth in the situations we've been considering in this chapter.[10] I should report that

---

[10] Throughout this chapter, I have been considering only cases where just two speakers are involved in a conversation, and have been ignoring questions raised by cases where there are more than two speakers. Suppose, for instance, that there are twenty people, say, in a meeting of a philosophy seminar, discussing whether a certain character 'knows' a certain fact. Suppose the seminar is not on the topic of skepticism, but rather concerns trying to formulate conditions for knowledge in light of Gettier

when I asked Stewart Cohen about the situations under consideration, he too was inclined to say that neither the skeptic nor her opponent was speaking the truth, so he is probably inclined toward something like the Gap view or the Exploding Scoreboard view.

Again, some may be worried about the extent to which the skeptic 'wins' on these views, since, though she is not construed as speaking the truth when she claims her opponent does not 'know', she does succeed at creating a context in which her opponents cannot truthfully claim that she does 'know'. But as I've indicated, I don't much care to what extent the skeptic 'wins' in the sense in question.

12. One-Way Disputes and the Asymmetrical Gap View

So, what of the cases (mentioned at the end of section 2) of 'one-way disputes', where a speaker, S1, has said, 'S knows', and a later speaker, S2, in a different conversation in which, on contextual analysis, higher standards seem to be operative, disputes S1's earlier claim: 'S1 was *so* wrong! S knows no such thing'? As I noted, the sense that S2's denial is incompatible with S1's affirmation can seem pretty strong here—almost as strong as in real two-way arguments, where both speakers are participating in the same discussion. But in one-way disputes, since the speakers are in different contexts, perhaps governed by different standards (on contextual analysis), it can seem that the 'single scoreboard' approach won't apply here, and the contextualist will have to somewhat counter-intuitively say about many such cases that

---

examples. Nineteen of the twenty converge on personally adopting a certain set of very ordinary standards, while just one personally indicates other, very different, much higher standards. The one hold-out does not play a leading role in the discussion, making soft-spoken contributions ('Well, I think Henry doesn't know, because for all he knows, he's a brain in a vat'), which are scoffed at and otherwise ignored by others, only every so often. With so much agreement on a particular set of standards in the room, are we to suppose that this lone hold-out creates a situation in which the 'know(s)' scoreboard explodes, or in which a huge gap is created in the truth-conditions of everyone's claims? I'm not sure what to say, but it's tempting here to resort to a Lewisism (from 'Scorekeeping', though Lewis uses it about a different kind of case) and say that if the character under discussion does meet the dominant standards in the room (the standards employed by the nineteen), then the claims of the nineteen that the character does 'know' are at least 'true enough.' Perhaps, though, the claims of the lone hold-out go gappy: When *he* claims that the character doesn't 'know', his claim is true iff the character doesn't meet either his own personally indicated standards *or* the dominant standards in the room; his claim is false iff the character meets both sets of standards; and if, as I'm supposing, the character meets the dominant standards in the conversation, but fails to meet our hold-out's personally indicated standards, the hold-out's claim is neither true nor false.

the denial and the affirmation are both true and are compatible with one another.[11]

First, note that this is a general problem that can arise for any context-sensitive terms, even obviously context-sensitive ones. Suppose the earlier speaker, S3, says 'Frank is here', in a context in which 'here' is being used to designate the city of Atlanta, where the discussion is taking place. Later, in a different discussion, S4 says, 'S3 was so wrong! Frank wasn't here at all.' Here, S4 seems to be vehemently disagreeing with S3. But suppose that in S4's conversation, the extent of 'here' that the speakers have been indicating covers just the hotel in Atlanta where these discussions are taking place. Well, then, should we rule that S3's affirmation and S4's denial are compatible, since it's certainly possible that Frank should be in Atlanta, but not in that hotel? But S4 seems to be quite clearly indicating that she means to be saying something incompatible with S3's affirmation! What to do? Well, that's not the easiest of questions. But, presumably, the solution *isn't* to give up on 'contextualism' and instead go 'invariantist' on the issue of how far out 'here' reaches in different contexts!

I'm inclined to take an 'Asymmetrical Gap' approach to these cases of one-way disputes. As we've just seen, my preferred treatment of real, two-way arguments is the Gap view described in the previous section. This Gap approach to two-way disputes seems the best way to take account of each speaker's 'personally indicated content', and also each speaker's indication that they mean to be saying things incompatible with what the other speaker says.

The Asymmetrical Gap approach that I'm now suggesting for cases of one-way disputes is similar, but posits a gap only in the truth-conditions of the later speaker's claim. The later speaker, after all, is indicating that she means to be disagreeing with the claim of the first speaker, but may also be in a conversation in which she and those she's talking with have indicated a certain content that differs from the content that governed the first speaker's claim. So, taking both of these indications seriously, the asymmetrical gap approach to such one-way disputes rules that, in the case above, S4's assertion 'Frank wasn't here at all' is true iff Frank was not even in Atlanta, false iff Frank was in the hotel, and truth-value-less if Frank was in Atlanta but not in the hotel in question. Similarly, S2's claim, 'Frank knows no such thing', said in dispute of S1's earlier 'Frank knows', is true iff Frank meets neither the epistemic

---

[11] Thus, John MacFarlane objects, 'The problem with [DeRose's] single scoreboard approach is that it explains only intraconversational disagreement, leaving inter-conversational disagreement unaccounted for' (2007: 21).

standards of S1's context nor those indicated by S2's later context, is false iff Frank meets both sets of standards, and is truth-value-less iff Frank meets one, but not the other, standard. But the earlier claims (by S1 and S3) are not rendered gappy by what happens in some later conversation.

This Asymmetrical Gap approach does not deliver the same, neat result that the simple Gap approach does that what the one speaker is denying is the very same gappy thing that the other affirms. But that seems appropriate to the cases—those of 'one-way disputes'—to which it is applied. And the Asymmetrical Gap view *does* deliver the result that the relevant denials are incompatible with the relevant affirmations: No matter where Frank is, there's no way that S3's affirmation and S4's denial can both be true; and no matter how well S is positioned with respect to p, there's no way that S1's attribution of knowledge and S2's denial of knowledge can both be true.

13. The Asymmetrical Gap View Applied to Relations between Earlier and Later Claims Made during the Same Two-Way Dispute

The Asymmetrical Gap view can also be applied to some claims made in two-way disagreements, and thereby solve a perplexity that arises about such cases. As we've seen, I apply the original Gap view (I guess we can now also call that the 'Symmetrical Gap' view) to two-way disagreements, after, as I put it, the debate has reached a 'state of equilibrium'. But what of the initial assertions made in such debates, before it is determined that the original assertion would be disputed by another speaker indicating different standards, and their relation to later claims that are intended to contradict them? The Asymmetrical Gap view I've just suggested for handling cases of one-way disputes can also be usefully applied to handle some of these tricky issues that arise within single conversations. Later claims and the earlier claims they dispute, made within a single, extended argument, can be treated in the Asymmetrical Gap way that I suggest for later claims and earlier claims in cases of one-way disputes that happen in different conversations. If conversational moves have been made to alter the epistemic standards between the times of the earlier claims and the later claims within a single conversation, yet the makers of those later claims indicate not only that they intend the new standards, but also that they are contradicting (or agreeing with, for that matter) the old claims, the truth-conditions of these new claims 'go gappy' between those of the old claims, and what the new truth-conditions would be were it not for the indications that

the new claims contradict (or agree with) the old ones. This way, the content of the earlier claims is kept appropriately independent from what transpires later in the conversation, while the incompatibility of the later claims with the earlier claims they are meant to oppose is preserved. Again, it's the conflicting intentions and/or indications of the relevant speaker—in this case, the one making the later claim—that underwrite the gappy truth-conditions that are assigned to her assertion.

14. Is There a Good Objection to Contextualism to be Found in its Inability to Handle Cases of Disagreement?

Throughout this chapter, I have stressed that, though we are especially interested here in the semantics of 'know(s)', the questions we are addressing arise also with other context-sensitive terms, and are in that way quite general. I will close with a brief discussion of what that means for the potential success of objections to contextualism based on its alleged inability to handle situations of disagreement.

About cases of face-to-face disputes, I have here indicated my preference for a Single Scoreboard Semantics approach, and, more particularly, for the Gap view. And in cases of what I'm calling 'one-way disputes', as I've just intimated, my own inclination is to adopt the Asymmetrical Gap view. Those are not only my inclinations about what to say about cases involving 'know(s)', but also for what to say quite generally about the truth-conditions of sentences containing context-sensitive terms in the relevant situations—as I've tried to indicate at a few places by explaining that I would give a similar treatment to the issue of the extent of 'here' in analogous conversational situations.

However, not everyone will be inclined to give the same answers as me to the general questions of what to say about the truth-conditions of claims involving context-sensitive terms in the relevant situations, and though I've sought to explain my reasons for preferring the answers I've chosen, I realize that I haven't presented anything close to a conclusive case for taking the particular path I've chosen.

I urge those who prefer a different approach to the general questions not to judge contextualism about knowledge by plugging my preferences on these issues into the contextualist position, thinking of the result as 'the contextualist position', and then counting it against epistemic contextualism that it doesn't handle situations of disagreement in a way you find attractive. That type of

reasoning can provide you with a reason for rejecting the particular version of contextualism that results from plugging in my own preferences on these issues, but not for rejecting contextualism generally. To evaluate epistemic contextualism in light of these potential objections concerning its handling of the relevant situations, you do better to take your own preferences on the general questions and see how contextualism about 'know(s)' looks to you in the relevant situations. Tell me how you are inclined to handle analogous issues about, say, how far out 'here' reaches, and I'll tell you how to handle those issues with regard to 'know(s)'.

Are there any *special* problems that arise for contextualism about 'know(s)' in these situations? If not, there doesn't seem to be much of an objection to epistemic contextualism to be found here. These are general problems for what to say about semantically flexible terms in certain situations. Presumably, there is some correct solution to these problems—unless you want to deny that there are any context-sensitive terms. Unless there's some special problem that arises for thinking that the correct solution, whatever it is, can be adopted by contextualists about 'know(s)', we haven't uncovered any strong objection to epistemic contextualism in these situations. What we have instead uncovered are some interesting general questions about context-sensitive terms and some of the options for how to deal with these questions.

5

'Bamboozled by our Own Words'

Semantic Blindness and Some Objections
to Contextualism

1. Methodology, Straightforward Data, and Objections to Contextualism Based on Fancier Features of Ordinary Usage

A certain methodology would strongly favor contextualism.[1] This 'methodology of the straightforward', as we may call it, takes very seriously the simple positive and negative claims speakers make utilizing the piece of language being studied, and puts a very high priority on making those natural and appropriate straightforward uses come out true, at least when that use is not based on some false belief the speaker has about some underlying matter of fact. Relatively little emphasis is then put on somewhat more complex matters, like what metalinguistic claims speakers will make and how they tend to judge how the content of one claim compares with another (e.g. whether one claim contradicts another). This methodology favors a contextualist conclusion, because, as I have been urging here, especially in Chapters 2 and 3, the data concerning what simple positive and negative claims we make involving 'know(s)' do strongly support contextualism. Meanwhile, considerations coming from the other, fancier sources can seem less kind to contextualism. One might therefore be tempted to read my defense of contextualism as an instance of someone following a certain contextualist-friendly methodology to its inevitable contextualist conclusion.

Now, I must admit, I do find the methodology of the straightforward somewhat attractive, at least when its favoring of the simple data is not taken too far. This methodology can be based on the thought that what speakers know best

---

[1] My thinking here was very much helped by Timothy Williamson's presentation at a conference on contextualism at the University of Stirling in March 2004—though Williamson's views on the matter are very different from my own. See Williamson (2005b: esp. 222–3). I was also helped here by a post Brian Weatherson had on his weblog, 'Thoughts, Arguments, and Rants', in which he complained that contextualists have focused too much on what I'm here calling the straightforward considerations.

about the piece of language at issue is how and when to make simple positive and negative applications of the term in question, and it does seem sensible to me to suppose that that is what speakers would be best at, at least to the extent that it should count as a somewhat worse strike against a theory that it runs afoul of the straightforward, as opposed to the more complex, data. And I believe that this sensible thought is often verified by experience—experience of the fancier data misbehaving in various ways. Thus, if push were to come to shove, I suppose I would be somewhat inclined to follow the data concerning the straightforward.

But I hold that we can forget about any such general biases here, however sensible they might or might not be, because I think the fancier considerations at least misbehave (and perhaps even work a bit against the invariantist) *in the case at hand*. Thus, even if we don't relegate such considerations to some secondary, suspect status, they won't speak clearly against contextualism, anyway. So push does not come to shove here.

In this chapter, then, I play defense, responding to several related objections to contextualism. First, an objection from comparative judgements of content: It is alleged that, where we again use 'HIGH' to designate the contextualist's high-standards case and 'LOW' to name his low-standards case, the denial of knowledge in HIGH seems to contradict the affirmation of knowledge in LOW, and that this is a problem for contextualism. Second, an objection from metalinguistic claims: It is alleged that if S-HIGH (the speaker in HIGH) and S-LOW (the speaker in LOW) were told of each other, S-HIGH would say, 'S-LOW's claim is (probably) false'[2] and/or that S-LOW would say 'S-HIGH's claim is false,' and that this is a problem for contextualism. Third, objections from belief reports: It is alleged that if S-HIGH and S-LOW were told of each other's claims, and each thought that the other's claim was sincere, S-HIGH wouldn't hesitate to say 'S-LOW believes that S knows' and/or S-LOW wouldn't hesitate to say 'S-HIGH believes that S doesn't know', and that this is a problem for contextualism. I will also seek to pre-empt a variety of objections based on how 'knows' behaves inside of both belief reports and speech reports by showing that in such contexts 'knows' behaves markedly like other terms that are clearly context-sensitive.

Objections of the types described in the above paragraph are sometimes accompanied by the charge that the contextualist's best response to them—though it is costly to have to reply this way—is to appeal to 'semantic blindness': to admit that in the ways indicated in the objection, speakers do talk

---

[2] 'Probably' may have to be inserted here because S-HIGH, being in a context where the standards for knowledge are unusually high, may also be in a context where the standards for flat-out assertion are unusually high, and thus may be in no position to flat-out assert much of anything, including that S-LOW's claim is false.

as if contextualism were false, but to claim that speakers do so because they are unaware of—blind to—the context-sensitivity of 'know(s)'. So, we will also look at a related objection Stephen Schiffer raised against contextualism (and in particular to the application of contextualism to the problem of skepticism), which seems to have inspired this talk of semantic blindness. I will argue that none of the objections succeeds, and that the way that the contextualist does implicate speakers in semantic blindness does not hurt the case for contextualism, because speakers are involved in equally problematic semantic blindness whether or not contextualism is true.

2. The Objection from Judgements of Comparative Content

S-LOW says that someone 'knows' something to be the case, while S-HIGH says that very same person 'doesn't know' that very same thing to be the case. It is alleged that these two seem to be contradicting one another, while contextualism counter-intuitively rules that the two assertions are perfectly compatible. It can appear that the contextualist's best hope here is to—quite lamely—appeal to semantic blindness. The two assertions are 'surface-contradictory', just as 'I am hungry' and 'I am not hungry', said by two different speakers, are surface-contradictory. Now, this latter pair don't even seem to us to be incompatible with one another, because the relevant 'context-sensitivity'—that 'I' refers to different people in the different claims—is quite clear to us. It can seem that the contextualist's best hope here is to claim that S-HIGH's assertion seems incompatible with S-LOW's because the two assertions are surface-contradictory, and, unlike what happens with 'I', we are blind to the relevant context-sensitivity that renders these two surface-contradictory claims compatible with one another.

But is the contextualist's verdict here really counter-intuitive? Not if we use the right test cases.

What will the contextualist's pair of cases—HIGH and LOW—be like? As we discussed in Chapter 4, many immediately imagine a dispute between two speakers, one of whom is insisting she does 'know' something and the other of whom insists that the first does not 'know' any such thing. And, more particularly, many immediately imagine just such a dispute between a philosophical skeptic and some Moore-like resister of skepticism—perhaps even Moore himself. In such cases of dispute, the intuition that the two speakers are contradicting one another can indeed be very powerful. (Correlatively,

the intuitions that S-LOW's claim to 'knowledge' and S-HIGH's denial of 'knowledge' are both true can be extremely weak to non-existent.) If the contextualist held that the claims of these two disputants were compatible with one another, this really would be a strong strike against contextualism—perhaps strong enough to compete with the strong grounds we saw in favor of contextualism in Chapters 2 and 3. But as we saw in Chapter 4, this isn't '*the* contextualist position'. It is one *possible* contextualist position. But contextualists don't have to take that position, and actual contextualists tend not to take it. I don't take that position, and if any actual contextualists do take it, I'm unaware of that fact.

In assuming that contextualists will judge both assertions in such a dispute to be true, it is assumed that contextualists will take what we have called (in Chapter 4) a 'Multiple Scoreboards' approach to such disputes, on which the content of each speaker's use of 'know(s)' follows that speaker's own 'personally indicated content'. But I, and, so far as I can tell, the other contextualists I know of, all opt instead for some version of the 'Single Scoreboard' approach, on which the content of both speakers' uses of 'know(s)' is the same, so that what the one speaker is affirming is exactly what the other is denying, and their assertions are not at all compatible. (Indeed, as we saw in Chapter 4, on the version of the Single Scoreboard approach I prefer, the 'Gap view', so far from its being the case that both of our disputants are asserting truths, neither of them is!)

So, if you claim that the parties in such a dispute are contradicting one another, don't be surprised to find contextualists agreeing with you. You're going after a straw contextualist. We need to instead check the right cases—cases about which the contextualist really will claim that both S-LOW's affirmation of and S-HIGH's denial of 'knowledge' are true.

Well, what of cases of what we might call 'one-way disputes', where a speaker, S1, has said, 'S knows', and a later speaker, S2, who has heard about what S1 said, and who is in a different conversation in which, on contextual analysis, higher standards seem to be operative, disputes S1's earlier claim: 'S1 was *so* wrong! S knows no such thing!'? The sense that S2's denial is incompatible with S1's affirmation can seem pretty strong here, too—almost as strong as in two-way arguments. But since these speakers are in different contexts, which can seem to be governed by different standards (on contextual analysis), it can seem that the 'Single Scoreboard' approach won't apply here, and the contextualist will have to somewhat counter-intuitively say about many such cases that the denial and the affirmation are both true and are compatible with one another. But, as we again saw in Chapter 4, contextualists easily can, and I, for one, do, adopt a plausible approach to such cases on which S2's claim is quite inconsistent with S1's.

So, what *are* the 'right' cases—cases about which the contextualist really will say that the affirmation in LOW and the denial in HIGH are both true—like? In Chapter 2, I described the best pairs of cases for contextualism. The contextualist certainly will hold that the affirmation of knowledge in LOW and the denial of knowledge in HIGH are both true concerning these best case pairs, and consequently that these two assertions are compatible with one another; after all, the contextualist's chief positive argument is based on just those verdicts. I described these best pairs of cases in some length in Chapter 2; I'll quickly summarize here. First of all, S-HIGH and S-LOW are not having any argument at all; indeed, they're not even talking to each other. They're in completely different conversations.[3] The speakers in LOW are applying relatively low epistemic standards to subjects, describing them as 'knowing' various things so long as they meet rather moderate standards. And such use of low standards is quite appropriate given the practical interests in play in LOW. But in HIGH, different interests are in play: for instance, the speakers may be facing a very high-stakes decision which makes the use of high epistemic standards appropriate, and the speakers in HIGH have indeed been applying such elevated standards, describing subjects as 'knowing' only if they meet unusually elevated standards, and going so far as describing subjects as not 'knowing' if they meet only the moderate standards used in LOW, but not the elevated standards of HIGH. However, the standards being applied in HIGH, while they are unusually heightened, still are not the ethereal standards that can seem to be in play in discussions of philosophical skepticism, and do seem to be appropriate to use given the heightened practical concerns in play in HIGH. In neither LOW nor HIGH is there any disagreement among any speakers about whether any subject 'knows'. When S-HIGH says that the subject does not 'know', she meets with no resistance at all from their conversational partners, who have also sensed that unusually stringent standards are appropriate in the context, and have adjusted their use of 'know(s)' to these high standards, themselves saying, with apparent truth and appropriateness, that they and other subjects do not 'know' things if those subjects fall short of the unusually high standards in question, even if they meet more ordinary standards. This best kind of high-standards case involves no element of arm-twisting, arguing, or convincing, but only of *informing*—one speaker informing others that, according to the appropriately heightened standards that obviously govern their conversation, some subject does not 'know' something. Paired off against such high-standards cases, it's also a feature of our use of 'know(s)' that

---

[3] S-HIGH and S-LOW may be, and indeed often are, one and the same person, imagined as being in two quite different conversations.

often, where unusually heightened practical concerns are not involved, speakers behave in ways that show that they are in a low-standards context, informing each other that the same subjects described above, or similarly positioned subjects, *do* 'know' those same things that they alleged not to 'know' in the high-standards cases. So it is in LOW. And where there is no dispute among the parties to the conversation in either of the cases; and where none of the speakers in either case is made to reverse herself, saying that something is not 'known' that she had earlier claimed was 'known', or vice versa; and where the different standards that seem to govern the two cases both seem appropriate, given the practical situations involved in each case, the intuitions that the relevant claims in each of the cases (the ascription of knowledge in LOW and the denial of knowledge in HIGH) are true can become extremely powerful indeed. (And the contextualist needs only some case pairs to work to win the day, while the invariantist needs all cases to be governed by the same, invariant epistemic standards.)

What becomes of the intuition that S-HIGH's assertion contradicts S-LOW's when we construct HIGH and LOW in the right way (so that contextualists really will say both assertions are true)? It is putting the matter extremely mildly to say that the intuition is then considerably weakened.

If you teach a class (perhaps in epistemology or philosophy of language) in which the issue is relevant, you can try the issue out on your students: Present to them *good* LOW and HIGH cases (but don't label them as 'LOW' and 'HIGH', since that might start biasing matters), and see whether they think there is a contradiction here. When seeking intuitions, it's often important that issues be presented as fairly and as neutrally as possible, because, as we all know, the manner of presentation can greatly affect the responses one receives. But I think you can forget all that in this case. Present the issue in a way that creates a strong bias against the contextualist answer. Indeed, rather than asking in any fair way whether there's a contradiction here, you can authoritatively assert that there is one, and then ask for feedback on a related question. Try this, for instance: 'Sue says that John knows; Mary says that he doesn't know. *Obviously*, these assertions contradict one another, so they can't both be true. So, I ask you: Which speaker is making a true claim, Sue or Mary?' If my past experience is any guide at all, in almost every class of sufficient size, *if you are using the right kinds of test cases* (and that qualification is vital), you will find that some student will overcome the strong anti-contextualist bias of the presentation, insist that both speakers are asserting truths, and then, if you're at all lucky, will even present a proto-contextualist analysis of the situation to explain how this can be. And once this happens, a significant proportion of this brave student's classmates will join her in her verdict that both are asserting

truths (though some squabbles may well break out within this camp over the details of the proto-contextualist analysis they're using to explain the situation). In fact, if you really do a good job of constructing LOW and HIGH, it's a pretty good bet that significantly more students will follow this contextualist suggestion than will stick with you on the issue of contradiction. But I'm on defense here, and am not looking for any positive support for contextualism. For my current defensive purposes, it's more than sufficient that there is no intuitive winner here. And that, at least, seems quite safe.

3. 'Semantic Blindness': Get Used to It!

Of course, you will get results pushing even more strongly in a pro-'contextualist' direction if you present speakers with a similar question involving a term that is more obviously context-sensitive than 'know(s)' is, like 'here':

> It's snowing in Chicago, but not in Los Angeles. Cher is in Chicago, talking on the phone with Lonnie, who is in Los Angeles. Cher says, 'It's snowing here.' Lonnie replies, 'It's not snowing here.' Obviously, these two assertions contradict one another, so they can't both be true. So I ask you: Which speaker is making a true claim, Cher or Lonnie?

Hopefully, *all* your students will go 'contextualist' here on 'here', insisting that both speakers are making true claims and are not contradicting one another. Contextualism about 'know(s)' doesn't fare nearly as well, I freely admit. So I suppose there must be something to the charge that if contextualism is true, we suffer from some degree of 'semantic blindness': Speakers are to some extent blind to the context-sensitivity of 'know(s)'.

But this is not a consideration that favors invariantism over contextualism, for speakers seem about equally afflicted by semantic blindness whether contextualism or invariantism is correct. Indeed, if contextualism is true, then you and the students in your class, if any, who stick with you on the issue of contradiction are 'blind' to the context-sensitivity of this very common verb, and I am sorry to have to charge you with such blindness. But, equally, if invariantism is instead correct, then the brave student who stands up to you in class, and all her classmates who follow her—and me, too, for that matter—are all 'blind' to the context-*in*sensitivity of that same common verb. To mimic the rhetoric one can often hear coming from critics of contextualism, and turn it against them: 'Are we *really* to believe that perfectly ordinary and competent speakers are *that* blind to the workings of their very own common word?!!' Well, yes, I guess we're stuck with that, either way. The intuitive situation

could only appear to be so strongly favoring invariantism because people were imagining the wrong kinds of cases—cases about which contextualists themselves can and often do deny that both speakers are asserting truths. Once we look at the right cases, this issue of semantic blindness appears to cut no ice at all. Look, when HIGH and LOW are constructed right, whether S-HIGH'S claim contradicts S-LOW's is just a tough question, about which there is no intuitively clear answer. So there *simply is* a good deal of 'semantic blindness' afflicting speakers here, whichever answer is correct.

4. The Objection from Metalinguistic Claims

The situation is similar enough with the objection from metalinguistic claims that I will discuss that objection only quickly. In this objection, recall, it is alleged that if S-HIGH (the speaker in HIGH) and S-LOW (the speaker in LOW) were normal speakers and were told of each other, S-HIGH wouldn't hesitate to say, 'S-LOW's claim is (probably) false', and S-LOW would quickly say, 'S-HIGH's claim is false', and that this is a problem for contextualism. And, again, it can appear that this is a problem for which the best solution is to lamely appeal to semantic blindness: The speakers would make these metalinguistic statements because they are blind to the context-sensitivity of 'know(s)', and so each thinks that the truth of their own assertion implies that the other's claim is false.

But again, it seems that the intuitions appealed to here are at all strong only if we are imagining the *wrong* kinds of test cases. If S-HIGH and S-LOW are having an argument with one another, then, yes, each will be inclined to assert the other's claim is false almost to the same extent that each is willing to make their own object-level claim. But when HIGH and LOW are constructed correctly—and we won't again go into what's involved in doing that—things are very different.

Perhaps it's best to briefly compare the intuitive situation here with the positive basis for contextualism—that S-HIGH's denial of 'knowledge' is true and that S-LOW's affirmation of 'knowledge' is true. As I've reported, when HIGH and LOW are constructed correctly—and especially when the standards employed in each case are appropriate given the practical situation being faced in each case—the intuitions that the assertions are true can get very powerful indeed, especially when each case is considered individually. In addition, as we can sense with perhaps even more confidence than that with which we directly judge the truth-values of the assertions, each assertion clearly seems

to be the appropriate thing for a speaker to say in the situation at hand. And, as I argued in Chapter 2, this supports the intuitions that the assertions are true, because there is a strong presumption that perfectly appropriate assertions are true, at least when they are not based on some false belief that the speaker has about some underlying matter of fact relevant to the assertion at hand.

By contrast, when HIGH and LOW are constructed correctly, the intuition that S-HIGH would be speaking truthfully if she said 'S-LOW's claim is false' is very weak indeed. In fact, I myself have always been fairly strongly inclined to judge that S-HIGH would then be saying something false, but that's me, and it is perhaps not a coincidence that I became a contextualist. But on this question, you will find a lot of speakers siding with me, *if you have constructed the cases well*. Yet beyond the question of the truth-value of such a metalinguistic claim, even the matter of whether such a claim is appropriate is quite vexed. About the right test cases, it has always quite strongly seemed to me that either speaker would be going too far to claim that the other's assertion is false, and again, many speakers would side with me on this.

Does LOW's claim contradict HIGH's? Should they each say that the other's claim is (probably) false? Would such metalinguistic claims be true? As I've noted, I've always been fairly strongly inclined to give the contextualist's answers (no, no, no) to these questions, at least when—and I can't emphasize this important qualification enough—they're asked about properly constructed test cases. But I have always taken these to be rightly disputed matters. Those are precisely the questions we need to do some philosophy to answer. I have a hard time relating to an enquiry that takes certain answers to *these* questions as its starting point, or that makes delivering certain answers to these questions a desideratum of a theory—especially when the answers being sought are to my mind the *counter*-intuitive ones! And, of course, to the extent that it is some slight advantage to a theory that it deliver intuitively correct answers to these questions, given my intuitive reactions to the questions, it's the theory that delivers the contextualist answers that gains that slight advantage, at least in my book—and also for the many speakers who have intuitive reactions similar to mine.

5. Hawthorne and Belief Reports

We turn next to belief reports—though similar objections could be made from other reports of propositional attitudes. If one neglects to consider the best

cases for testing such a general principle, it can seem that our use of 'know(s)' within belief reports quite generally follows a 'disquotational schema' like this:

> (DSK) If a speaker S1 sincerely utters a sentence of the form 'A knows that *p*', where 'A' is a name that refers to a subject, then a speaker S2, who is using 'A' with the same meaning as S1 and to refer to the same subject S1 refers to, and who is using 'p' with the same meaning as S1 and to designate the same proposition that S1 designates, can truthfully assert a sentence of the form 'S1 believes that A knows that p', where 'S1' is a name that refers to S1.

I here formulate the principle to apply to cases where *names* are used to refer to the people involved. The principle could be formulated so as to apply as well to cases where other devices—like pronouns—are used to designate the people involved, and so could be applied to cases like 'He believes that he knows that he won't die of a heart attack.' However, this would make the principle quite complicated. So in discussing cases, I'll just assume that we understand how such a principle would apply to examples involving such devices, without explicitly formulating the principle so it can apply to such cases. This simplifying move should not harm our discussion, since in my attack on DSK, I won't pursue criticisms whose fates depend on fine points of the formulation of the principle, but rather attack its basic idea as it applies to cases where the alleged context-sensitivity of 'know(s)' is at issue. The above formulation doesn't have the basic structure typically associated with what are called 'disquotational' schemas, so some readers may wish not to call it such a schema.[4] *I* call it a disquotational schema for the following reason. If S1's utterance of 'A knows that p' is sincere, then, even if S2 is in a very different context, she can of course truthfully assert, 'S1 believes what she asserts by the sentence, "A knows that p".' But that is *so* clunky! What DSK tells us is that S2 can (also) come down a level ('disquote') and truthfully state the far simpler, though still a *bit* clunky, 'S1 believes that A knows that p.'

It can seem that if 'knows' were context-sensitive, then no disquotational schema like this would govern its use, because if the epistemic standards that govern S2's context are different from (and especially if they are higher than) those that governed S1's context, S2 will seem to be ascribing to S1 a

---

[4] 'Disquotation'/'disquotational' principles or schemas are usually presented so that their consequents contain a simple, object-level use of the term in question; for example:

If S sincerely asserts, 'A knows that p', then S believes that A knows that p.

Hawthorne's own formulation of DSK (discussed below in n. 7) is in that way more in keeping with typical disquotational schemas than is the formulation I am using.

belief—that A knows according to the higher standards—that S1 may well not hold.[5] Thus, it's no surprise that John Hawthorne is able to use DSK, together with other principles, to stir up trouble for contextualism (at 2004: 98−107). We will not look at the rest of Hawthorne's case, for it's at DSK that I will challenge him.[6] And Hawthorne thinks our ordinary usage does betray an *extremely* clear commitment to a fully general principle like DSK;[7] he writes:

It seems altogether obvious that we take sincere uses of the verb 'know' as a straightforward guide to the contents of what people believe. For example, if Jones says 'I know I have feet' and I reckon he is being sincere, I will have no hesitation at all in reporting the contents of his mind as follows:

(1) Jones believes that he knows that he has feet …

Of course, not all words behave disquotationally in this way. If a speaker says 'It is raining here', I cannot report the contents of his mind by an ascription of the form 'He believes that it is raining here', unless he and I happen to be in the same place. By contrast, it looks very much as if we *do* adopt something like DSK. If, for example, someone sincerely utters 'I know that I will never have a heart attack', we have no hesitation whatsoever in reporting the contents of his mind by claiming that he believes that he knows that he will never have a heart attack. That is how the verb 'know' seems to work.… DSK seems exactly right.   (2004: 100−1)

But so far from being 'altogether obvious' and 'exactly right', that just seems wrong, or at the *very* least, extremely unclear, *at least as applied to the best and most relevant test cases*. The best test cases for this general principle, DSK, are ones where the contextualist will think there is a marked difference in epistemic standards.

---

[5] Note that I'm only saying that this is how things 'can seem'. Whether DSK, if true, would really give the lie to contextualism depends on some very subtle issues especially concerning the semantics of speech reports.

[6] In one important respect, the basic structure of Hawthorne's argument at (2004: 98−107) is a bit opaque to me, in a way I explain in the appendix to this chapter. At least the main argument he intends clearly rests on DSK, and I will address that main argument here. But it's possible that at (2004: 104−7) he is offering an independent anti-contextualist argument that does not depend on DSK. I address that argument in the appendix to this chapter.

[7] Hawthorne's version of DSK is at (2004: 101). I based my formulation on Hawthorne's, but modified it by making its consequent metalinguistic, as seems appropriate to the issues at hand. In any case, the exact formulation of DSK should have no effect on my challenge to Hawthorne's argument, because I do not question whether the observations about ordinary usage on which he grounds DSK, if correct, would successfully support the principle (on that matter, the exact formulation may be very important), but rather argue that his observations are just not generally correct in the first place. What's important to our dispute, then, is that Hawthorne's DSK, like the variation I'm using here, is a general schema, which is supposed to cover all cases in which a speaker has issued a knowledge attribution, including cases where the contextualist would hold that there is a marked difference in epistemic standards between the context in which the belief report is made and the context in which the original speaker attributed 'knowledge'.

So, let's think about such cases. A slight variation of the cases we used in section 2 of Chapter 1 will serve us well here. So, as in Chapter 1, suppose that Thelma and Louise (we only need two characters here, so we're cutting Lena out of the story) have just been at the office where they work, where they encountered very strong evidence that their colleague John was at the office: They saw his hat hanging in the hall outside his personal office, as has always been a very reliable indicator of his presence, and they heard other colleagues talking to each other in a way that strongly indicated that John was there. Now, changing the story a bit, we'll suppose that both Thelma and Louise go to the local tavern, the setting for LOW, where they take part in a discussion with other colleagues concerning who has to pay up on a mere $2 bet that had been made ahead of time on whether the often-absent John would show up at the office today. As is appropriate to the context, speakers are describing subjects as knowing propositions if those subjects meet moderate epistemic standards. Just so, though they did not themselves see him at the office, Thelma and Louise are both claiming to 'know' that John was there, based on evidence that's described above, which seems more than sufficient for this discussion at the tavern. That evidence is accepted as quite sufficient for Thelma and Louise's listeners at the tavern, who are themselves at present disposed to describe subjects as being 'knowers' if they are in an epistemic position like that Thelma and Louise are in with respect to John's having been at the office.

Now we'll suppose that Thelma remains at the tavern, where she will continue to discuss the bet with more colleagues who are presently arriving. Louise, however, leaves the tavern to walk home, and soon finds herself in a very different conversational context, HIGH, where, at least on contextual analysis, much higher epistemic standards seem to prevail. Louise is stopped by the police, who are investigating some horrible crime, and, in connection with that, question her about whether John was at work that day. It emerges that they already have some reason to think he was at work, and no reason to think that's false (other than the fact that John is often absent from work, which Louise already knows), but, given the importance of the matter, they are seeking to verify that he was there. Our speakers in HIGH have executed conversational maneuvers for driving up the epistemic standards, and Louise is admitting that she does not 'know' various propositions if she is not in an extremely strong epistemic position with respect to them. Thus, even though Louise has grounds that would usually suffice for claiming to 'know' that John was at the office, and that actually did suffice for such a claim while she was at the tavern, she is in her present context admitting that she does 'not know'

that John was in—though, she adds, she has good reason to think he was in. In this context, she would claim to 'know' he was in only if, say, she has a clear recollection of having herself seen him there.

Now suppose the police, still interviewing Louise, ask whether Thelma might 'know' whether John was at the office. Louise, of course, knows that Thelma too heard the report and saw the hat, but was only at the office for a brief time and did not herself see John there. In short, as Louise knows, Thelma is in the same epistemic position with respect to John's having been at the office that Louise is in. How then will Louise, while in HIGH, describe Thelma? If Louise is a normal speaker, then even though she's aware of the fact that Thelma is presently in a very different type of conversational context, Louise will say that Thelma too 'doesn't know' whether John was in (though she may well add that Thelma, too, has good reason to think that John was in). After all, Thelma has no better reason for thinking John was in than Louise has, and Louise is admitting that she herself 'does not know'.

Now, for the question crucial to our current interests: While in HIGH (while talking with the police), will Louise say, 'Thelma believes that she knows that John was at the office yesterday'? As I've already intimated is my judgement about such claims, so far from being clearly correct, that claim seems to me at least somewhat clearly wrong. Yes, as Louise heard, Thelma said, 'I know John was at the office yesterday.' But Louise too made just such claims about herself while she was at the tavern. And now Louise is telling the police that she 'doesn't know'. And it seems quite wrong for Louise to say, while in HIGH, that either she or Thelma 'believes'/'believed' that she 'knew'/'knows'.

Yes, I know that there's all different kinds of wrongness we might be sensing here, and that consequently there's plenty of room for someone to maintain that such belief reports, while they're *somehow* wrong to make, would nonetheless be *true* if they were made. Maybe. Tough issue. But I'm playing defense here: This is the anti-contextualist's argument (at least as it's applied to the right kinds of test cases). And it is *far* from clear that such belief reports would be true, and far from clear that they would be appropriate. So any argument that starts from the intuitions that such reports are obviously true and/or appropriate is in deep trouble, unless it's accompanied by some good reason, not provided by Hawthorne, for why we should accept his general principle's verdicts in the best test cases for the principle, where there seems a marked difference in standards between the context of the belief reporter and the context of the speaker who attributed 'knowledge', and where, as it turns out, these verdicts seem not at all clearly correct.

Indeed, I suspect the reason Hawthorne so confidently writes that DSK seems 'exactly right', and claims that it's 'altogether obvious' that speakers show 'no hesitation whatsoever' in issuing the relevant belief reports in the relevant circumstances, must be that he never considered the best test cases for his principle. For how could one consider such cases and yet come up with such a confident verdict? When we think about the right test cases, more than just hesitation seems called for.

It's in response to the problem for contextualism that he derives from DSK (together with some other materials), that Hawthorne claims the contextualist is driven to appeal to semantic blindness:

> How should a contextualist in epistemology respond to all this? The most promising strategy would be to adopt a view that invokes *semantic blindness*: There is a real sense in which users of the word 'know' are blind to the semantic workings of their language. English speakers proceed as if DSK were correct. But it isn't. While those speakers are, at least implicitly, well aware of the context-dependence of 'tall', the same is not true for 'know'.   (2004: 107–8)

And Hawthorne rightly claims that having to resort to such an appeal is a drawback: 'Other things being equal, a semantic theory that didn't claim semantic blindness on the part of competent speakers would be a better theory' (2004: 109).

But, so far as this argument of Hawthorne's goes, the contextualist need not get to the point of having to make this problematic appeal. For the trouble Hawthorne stirs up is derived from DSK, which in turn is supported by appeals to how we report beliefs involving 'know(s)', and this argument should be stopped right at the outset, because, as we have seen, Hawthorne's general observation here is simply wrong. About the best test cases, where there is a marked difference in standards, it is very far indeed from being altogether obvious that we display no hesitation whatsoever in engaging in the relevant disquotation.

6. 'Know(s)' and 'Tall': A Better Objection Involving Belief Reports

The above only responds to Hawthorne's particular argument. Naturally, there are many potential arguments against contextualism by which it might be argued that within the likes of belief reports and other propositional attitude

reports, and also speech reports, 'know(s)' behaves in such a way as to cast doubt on contextualism, and at least one of these ways of attacking contextualism seems more promising than Hawthorne's objection. And naturally, I can't anticipate every possible way such an argument might go. Additionally, evaluating many such arguments would involve us in controversial general matters about how context-sensitive terms actually behave in such contexts, and settling such matters—as we are taught to quite conveniently write—goes beyond the scope of the present work!

But what I will do here in sections 6–9 is look at several different types of conversational situations, and compare how 'know(s)' behaves with how some clearly context-sensitive term behaves within belief reports and speech reports. By showing that in the relevant respects 'know(s)' behaves very much like the clearly context-sensitive term in the various situations, I will be able to answer the most effective way I have come up with of using our disquoting behavior in an argument against contextualism, and may also provide some grounds for thinking more generally that some powerful anti-contextualist argument in the vicinity will be hard to find.

In the above section, we tested DSK against cases where, at least on contextualist analysis, there was a very notable and obvious difference in epistemic standards between the contexts of our two speakers. As we saw, speakers really do, and should, become queasy about 'disquoting' in such cases of large and noticeable differences in standards.

But I admit that speakers are quite quick to engage in disquotation where whatever differences in epistemic standards there might be are small and not very noticeable. And, for a variety of possible reasons, this willingness on our part to so quickly engage in disquotation on 'know(s)' in such cases may be thought to spell trouble for contextualism—more serious trouble than is threatened by Hawthorne's argument. Here is perhaps the most plausible way to use this feature of the behavior of 'know(s)' against contextualism: There are many such cases where speakers breezily engage in disquotation where, it can be claimed, if contextualism is true, they should suspect that they might well be mis-reporting the content of a subject's belief. Example: I quickly walk past a group of speakers, and hear Frank say, 'Mary knows that the library closes at 5.' Later, I'm engaged in a different conversation, with a completely different group of speakers. But if the occasion calls for it, I won't hesitate to claim, 'Frank believes that Mary knows that the library closes at 5.' But, from my point of view, what are the chances, if contextualism is true, that the epistemic standards of my current conversation exactly match those of

the conversation I overheard? Quite small, it can seem. Perhaps that's not so much of a problem if the only worry is that the overheard conversation may have been governed by higher epistemic standards than is the conversation I'm currently engaged in, because if Frank believes Mary meets even higher epistemic standards, then he presumably believes she meets lower standards as well. But the case can be constructed so that, from my point of view, the standards governing the overheard conversation were slightly lower than those of my current context is about as likely as that they were higher. I just can't tell whether the standards that governed that overheard conversation were higher or lower than those of my current context, if, as can seem quite likely on contextual analysis, they differed at all from my current context. Not that I care much: As a normal speaker, I'm still pretty quick to disquote. And this can seem to be a problem for contextualism, because in saying the disquoted 'Frank believes that Mary knows that the library closes at 5', I would seem to be attributing to Frank the belief that Mary meets the standards of my current conversation, while my evidence for this claim—Frank's assertion, made in a different conversation which may well have been governed by lower standards—seems insufficient to adequately underwrite my seemingly confident, breezy attribution. Here, as in Hawthorne's argument, the critic appeals to our habits of disquoting to stir up trouble for contextualism. But the critic we're now considering doesn't use her observations to support some principle as general as DSK. Going through such a principle, we've seen, is a mistake, since such a principle applies to cases involving marked differences in standards, where, as we've seen, it falters. Rather, without trying to establish such a general principle, our current critic more wisely attempts to argue that some of the cases where disquotation *does* seem perfectly fine cast doubt on contextualism, even if there are also situations in which the relevant sort of disquotation seems problematic.

Ultimately, we should want a good account of belief reports, and how context-sensitive terms behave within them, that will explain why we're quick to disquote on 'know(s)' in cases like the one I've just described. Several candidate explanations suggest themselves, but each is too controversial to provide adequate relief from the attack. At this point, the issues of how belief reports work, and how context-sensitive terms behave within them, is too controversial for such a response. But we can effectively respond to this attack by noting that other terms that are clearly context-sensitive display relevantly similar behavior in analogous cases.

I assume that the gradable adjective 'tall' is context-sensitive: It's at least fairly obvious that how tall a person or an object must be for it to be true to

say of it that it is 'tall' is a context-variable matter.[8] Some may deny this,[9] but not many, and, to my thinking, not plausibly. And in situations analogous to the one described above, we're pretty quick to issue disquoting belief reports involving 'tall', too. Suppose what I hear Frank say as I quickly walk past is 'Mary is tall'. In a new conversation, I'll be pretty quick to say, 'Frank believes that Mary is tall', even though it seems I should consider it very possible that

[8] Some hold that gradable adjectives like 'tall' cannot provide a good model for the contextualist about 'know(s)', due to important differences in the behavior of the two terms. Stanley, in particular, has pressed this case forcefully at (Stanley 2004: 123–30, 134–9; updated at Stanley 2005: 35–46, 56–73). (Concerning a side point Stanley makes in this material, I here take the opportunity to register that I have a reaction to one of the examples Stanley uses that's very different from Stanley's own: 'It is more possible that Hannah will become a philosopher than it is that she will become a mathematician' seems far from unproblematic to me, though Stanley seems to think it's fine (Stanley 2004: 128). So in at least one respect, I think 'know(s)' may operate more like modal terms than Stanley thinks.) Peter Ludlow (2005: esp. 21–7) provides a good and interesting response to Stanley. Daniel Halliday (2007) interestingly proposes that the likes of 'sufficiently tall', which is also clearly context-sensitive, might make better models for 'know(s)' than do the likes of 'tall', and that at least most of Stanley's contrasts fail when the contextualist switches models in the way Halliday proposes. Stanley responds to Halliday—though I think unconvincingly—at Stanley (2005: 43–4). An important part of Stanley's case is his claims that the standards that govern gradable adjectives switch mid-sentence in examples like 'That butterfly is large, but that elephant is not large' (2004: 135; 2005: 58), which, it seems, can be truthfully asserted in some contexts, despite the fact that the elephant in question is larger than the butterfly in question, because the butterfly is large for a butterfly, while the elephant is not large for an elephant. However, at DeRose (2008: 155–8), I argue that such examples are best handled by holding, not that the standard governing the adjective changes mid-sentence, but that a single, 'group-relative' standard governs it throughout the conjunction. Another important difference alleged by Stanley is addressed in section 8 of this chapter.

In the final analysis, I suppose I do think gradable adjectives provide at least a *fairly* good model for how 'know(s)' operates—and a better model than many others seem to think, though this is based in part on my own understanding of the operation of gradable adjectives. (In addition to my handling of examples like 'That butterfly is large, but that elephant is not large' other conclusions I draw in DeRose 2008 are relevant here.) However, my current appeal to the behavior of gradable adjectives does not require the claim that these behave just like 'know(s)' does in all respects, or even that they provide a particularly good model for the overall behavior of 'know(s)'. Recall from section 8 of Chapter 2 that I am not—here nor anywhere in this volume—attempting an indirect positive argument for contextualism that takes as its premiss that some other term, like 'tall', is context-sensitive, and then argues that because 'know(s)' is so similar to 'tall', 'know(s)' too is context-sensitive. As I admitted, such an indirect argument would be very insecure, in my opinion, because, while there are very important similarities between the behaviors of 'tall' and 'know(s)', there are also many important differences. (We are, after all, comparing a verb with an adjective!) Here, I am only appealing to the behavior of 'tall' to rebut certain objections to contextualism about 'know(s)'. And the objections we will be considering here in sections 6–9 do seem to be soundly neutralized by the fact that the behavior of 'know(s)' that these objections are based on is also ways in which 'know(s)' behaves like the clearly context-sensitive term 'tall', even if there are also other ways in which the two terms behave quite differently. (I suppose that how well the behavior of 'know(s)' matches that of the likes of 'tall' in a wide variety of different respects would be very relevant to the type of objection to contextualism that I mention in the last two paragraphs of section 5 of Chapter 1. However, for reasons I give there, I can't come up with a version of that objection that is now very threatening, nor have I seen a threatening version of that objection.)

[9] See some recent work by Cappelen and Lepore, including especially their (2005).

the standards governing the true use of 'tall' in Frank's conversation were slightly lower than are those governing the new conversation I'm presently engaged in. So long as it appears that the standards of Frank's old conversation are even roughly in the same ballpark as are the standards governing my current conversation, I'm pretty quick to disquote—and here I think I'm behaving very much like a normal speaker. Why do we do this? I won't speculate here—though I do have some ideas, as do many of you, I suspect. But what's important to answering the charge against epistemic contextualism currently under consideration is that the behavior of 'know(s)' here is relevantly similar to that of the clearly context-sensitive term 'tall'.

As we noted in the previous section, when the epistemic standards governing the original speaker's knowledge attribution are clearly and markedly lower than are those that would govern the later attribution of belief, we do resist disquoting. (Recall that it seems wrong for Louise, while talking to the police, to say that 'Thelma believes that she knows that John was at the office'.) Of course, 'tall' behaves similarly. Suppose that while discussing the reasons why the 6 foot 5 inches Reeves is such a great high school basketball player, Lucius has sincerely said, 'Reeves is tall', and suppose that Theo knows that Lucius sincerely said that in such a context. Now, however, Theo is in a context in which various players, including Reeves, are being discussed as potential NBA centers, and, as is appropriate when potential NBA centers are being discussed, 6 foot 5 inches players are definitely *not* being counted as 'tall', and so Theo has admitted, 'Well, Reeves is not tall'. (Perhaps adding: 'But he has great low post moves, and he's young enough that he may still grow some more.') Here, as in the high-standards Thelma/Louise/'knows' case, it does not intuitively seem right for Theo to report in his 'high-standards' context, 'Lucius believes that Reeves is tall', for, though Lucius said, 'Reeves is tall', he said that in a different context where lower standards were in play. Again, there's room for one to claim that such a report, though quite inappropriate, would nonetheless be true. But there's room here to claim a lot of things. All I'm reporting here is that the behavior of 'tall' here is much like that of 'know(s)': the belief report is out of place, seems inappropriate and at best misleading, and normal, competent speakers wouldn't issue it.

7. 'Know(s)' and 'Tall': Some Speech Reports

If our speakers can't issue the disquoting belief reports we're considering, what about speech reports? Here it seems that 'know(s)' again follows 'tall' quite

closely. For with both terms, in cases of large, noticeable, and important shifts in standards, the relevant speech reports seem a bit, but only a bit, less wrong than do the belief reports we've already considered: It does seem a bit less wrong for Theo to report, 'Lucius said that Reeves is tall', and also for Thelma to report, 'Louise said that she knows'. But though these 'said that' reports seem less bad than the corresponding 'believes that' reports, they still seem at least somewhat problematic—at least a little hesitation seems called for. (Whether such 'said that' reports would be *true* is a very difficult question; they may just be in various ways misleading. I'm here merely reporting intuitions about whether these are the right things for our speakers to say.) But we've so far considered only the possibility of the speakers issuing the relevant speech reports and then shutting up. And what *does* seem clearly correct for the speakers to instead say—in both the case of 'know(s)' and of 'tall'—is the above 'said that' reports *with the addition of any of a number of possible follow-ups* that at least seem to work precisely by warning the listener against interpreting the content of what Lucius or Louise said as if it had been said in the listener's new, quite different context: 'Lucius said that Reeves is tall, but he was speaking of Reeves as a high school player,' 'Louise said that she knows, but she was speaking casually.' Something like that seems quite clearly to be the thing to say. Again, it's hard to resist trying to explain why this is so, but I will indeed resist. All I mean to be doing is pointing out the relevant behavior of 'know(s)' is quite similar in these contexts to that of the clearly context-sensitive 'tall'.

8. 'Know(s)' and 'Tall': 'I Never Said That!'

The last types of behavior we'll compare concern affirmations and denials by speakers that they had said something earlier in situations in which the speaker's consistency is being challenged because of the things she said in different contexts. Stanley thinks contextualism is in trouble because he thinks 'know(s)' here behaves differently from other context-sensitive expressions;[10] in fact, this is a major part of Stanley's case against contextualism. However, as we'll see, for contextualism there is really no trouble, but only good news, to be found here, because 'know(s)' actually behaves in a way very similar to the behavior of 'tall'—and similar points would hold for the comparison of

---

[10] (Stanley 2005: 51–4). Stanley also thinks that contextualism has trouble with cases in which a speaker uses 'what I said' to refer to an earlier claim of hers (cases of propositional anaphora) (2005: 54–5). I will deal with such phrases in Chapter 6, section 10.

the behavior of 'know(s)' with other gradable adjectives, and also with many other context-sensitive terms.[11]

As we've noted, while in HIGH (while talking to the police), Louise will say that she does 'not know' that John was at work that day, even though she has claimed to 'know' that very fact while she was in LOW (at the tavern). But now suppose the police officer questioning her has been in radio contact with another police officer who has arrived at the tavern since Louise left it, and challenges her as follows: 'Hey, but didn't you say at the tavern that you *did* know that John was at work today?' It isn't obvious how Louise might best answer, but one thing that is fairly clear here is that it would be wrong for her to reply, 'I didn't say that' or 'I never said that'. I should stress that by claiming that this response would be clearly wrong, I'm not asserting that it's clear that Louise would be saying something false if she were to say that. *No* verdict about the truth-value of such a problematic claim seems *clearly* correct. Again, I'm only reporting surface 'behavior' here. And though the exact nature of the wrongness involved in such response is not easy to diagnose, it is pretty clear that the response in question is *somehow* wrong. And, interestingly, Louise's response still seems wrong if she adds a clarifying explanation in the following way: 'I didn't say that. I was speaking casually then.' In line with what I noted in the previous paragraph, it seems much better (and, in fact, fine) for her to instead say, 'I did say that, but I was speaking casually then.' If she wants to give a fuller account that addresses how her earlier claim at the tavern and her later claim to the police were both somehow correct, there seems to be no entirely smooth way of putting that all together, but what seems acceptable is to attach clarifying devices to both halves of that in something like the following way (though a variety of clarifying devices would seem to do the trick): 'I did say that, but I was speaking casually then. I don't know that he was at work by the standards relevant to our current situation.'[12]

In the relevant analogous situation, 'tall' behaves in a remarkably similar way. While discussing Reeves's prospects as an NBA center, Theo will admit, 'Reeves is not tall.' But suppose the NBA scouts he is speaking with have heard reports about Theo's earlier conversation with his friends about what a great high school basketball player Reeves is, and challenge him as follows: 'Hey, but didn't you say to your friends that Reeves *is* tall?' As with the analogous situation with 'know(s)', while it isn't entirely obvious how Theo

---

[11] The context-sensitive term that Stanley compares with 'know(s)' is 'possible'. I address Stanley's comparison in n. 13, below.

[12] Another line of response that seems somewhat acceptable, at least to some (though others don't like it at all, I find), is for Louise to say: 'I don't know that John was at work today. I was speaking casually at the tavern when I said I did know it.'

might best answer, it is clearly wrong for him to reply, 'I didn't say that' or 'I never said that'.[13] And, continuing the parallel, Theo's response still seems wrong if he adds a clarifying explanation in the following way: 'I didn't say that. I was speaking of him as a high school player then.' The parallel grows even stronger as we next note that it seems much better (and, in fact, fine) for Theo to instead say, 'I did say that, but I was only talking about him as a high school player then.' And, finally, the parallel becomes almost eerie as we note that if Theo wants to give a fuller account that addresses how his earlier claim to his friends and his later claim to the scouts were both somehow correct, while there seem to be no entirely smooth ways of putting that all together, what seems acceptable is to attach clarifying devices to both halves of that in something like the following way (though, again, a variety of clarifying devices would seem to do the trick): 'I did say that, but I was speaking of him as a high school player then. He isn't tall [for an NBA center/by the standards relevant to our current concerns].'[14]

9. 'Know(s)' and 'Tall': Summary

To sum up the previous three sections, 'know(s)' tracks 'tall' impressively through the scenarios we've discussed here. In both the case of 'know(s)' and that of 'tall':

---

[13] As remarked in n. 11, the context-sensitive term that Stanley compares 'know(s)' with is 'possible'. He constructs the following dialogue:

A: It's possible to fly from London to New York City in 30 minutes.
B: That's absurd! No flights available to the public today would allow you to do that. It's not possible to fly from London to New York City in 30 minutes.
A: I didn't say it was. I wasn't talking about what's possible given what is available to the public, but rather what is possible given all existing technology.

Stanley declares that A's last line 'seems perfectly appropriate, and indeed true' (Stanley 2005: 53), which contrasts sharply with his verdict about the last line of an analogous dialogue concerning 'know(s)' (2005: 52). I strongly disagree. Whether A's last line in the above dialogue involving 'possible' (and, in particular, the first sentence of that last line) is *true* is a vexed question about which I have only a very tentative opinion, not worth sharing here. (And whether an analogous claim concerning 'know(s)' is true would be similarly vexed.) But it seems clear that that last line is *not* 'perfectly appropriate'; indeed, it seems not appropriate at all—and no more appropriate than an analogous claim involving 'know(s)'. The second sentence of A's last line is a perfectly fine clarification. But the first sentence seems a real clunker. There are many appropriate ways for A to respond which retain the second sentence. Some of them replace the first sentence with something quite different, but the easiest such response to describe simply omits that first sentence, making do with the second, clarifying sentence by itself.

[14] Another line of response that seems *somewhat* acceptable, at least to some (though others don't like it at all, I find), is for Theo to say: 'Reeves isn't tall. I was only speaking of him as a high school player when I said that he was tall.'

(a) The relevant disquoting belief reports seem fine where there's no clear and marked difference in standards between the context of the report and the context of the assertion on which the report is based.

(b) The relevant disquoting belief reports do seem somehow wrong where there is such a clear and marked difference in standards.

(c) The analogous disquoting speech reports also seem somehow wrong—though perhaps a bit less so than the belief reports—where there is a clear and marked difference in standards, and the speech report contains no clarifying follow-up.

(d) The analogous disquoting speech reports seem perfectly fine—in fact, they seem exactly the things to say—even where there is a clear and marked difference in standards, when they do include any of a number of clarifying follow-ups.

(e) In the relevant situations of challenge of the type we've discussed, after a change of context, it doesn't seem right to defend your earlier statement by insisting, 'I didn't say that', and it also seems wrong to reply by saying, 'I didn't say that', following that claim up with a clarifying remark about the context in which the earlier statement was issued.

(f) In the relevant situations of challenge of the type we've discussed, after a change of context, it does seem fine to admit, 'I did say that, but…', following up with a clarifying remark about the context in which the earlier claim was made. It also seems acceptable to then follow all of that up with a restatement of the later claim made in conjunction with a clarifying device concerning it.[15]

Thus, it's hard to see how any of this behavior of 'know(s)' can serve as a good foundation for an argument against the context-sensitivity of 'know(s)', given these similarities in the behavior of the clearly context-sensitive 'tall'.

10. Schiffer's Attack on Contextualist Solutions to Skepticism: Being 'Bamboozled by our Own Words'

In a note attached to the passage I quote toward the end of section 5, above, where he introduces the notion of 'semantic blindness', Hawthorne suggests

---

[15] Alternatively, it seems somewhat acceptable, at least to some (though other people don't like this response at all), to instead reply to the challenge by reissuing the later claim without clarification in conjunction with a clarifying explanation of the earlier claim: see nn. 12 and 14, above.

that we see Stephen Schiffer's (1996) 'for related discussion'. I will close this chapter by following that excellent suggestion.

Schiffer's attack is directed specifically at one of the most important proposed applications of contextualism—to the problem of philosophical skepticism. And while Schiffer doesn't (at least so far as I can find) use the phrase 'semantic blindness', he does seem to have the same idea that Hawthorne designates by that phrase when, for instance, Schiffer issues the memorable charge that the contextualist solution to the problem of skepticism essentially involves 'a pretty lame account of how, according to [the contextualist], we came to be bamboozled by our own words' (1996: 329). At any rate, it seems that Schiffer's attack can be effectively answered in a way very similar to how I respond to charges of 'semantic blindness' above in section 3.

Schiffer's attack is directed at the contextualist response to this skeptical argument:

[SA] 1. I don't know that I'm not a BIV (i.e., a bodiless brain in a vat who has been caused to have just those sensory experiences I've had)

2. If I don't know that I'm not a BIV, then I don't know that I have hands

So, C. I don't know that I have hands. (317)

As Schiffer understands the contextualist's account of [SA]'s persuasive power, the skeptic, in asserting her first premiss, raises the standards for knowledge to a level (which Schiffer labels the 'Tough' standards) at which one counts as knowing neither that one is not a BIV nor that one has hands. Thus, contextualists, as Schiffer understands them, hold that when the skeptic presents it, the sentences of [SA] express this argument:

[SA-T] 1. I don't know that I'm not a BIV relative to Tough

2. If I don't know that I'm not a BIV relative to Tough, then I don't know that I have hands relative to Tough

So, C. I don't know that I have hands relative to Tough (324)

And [SA-T], which is the argument that [SA] expresses, is not only valid, but is also sound. Why, then, does the argument strike us as at all paradoxical? Why don't we simply accept this sound argument that the skeptic presents? Schiffer continues:

Naturally, the Contextualist realizes that he isn't quite finished. He must explain why [SA] *seemed* to present a paradox. If the argument [SA] *really* expresses is plainly sound, then why do we instinctively feel that [SA] expresses an argument that's

plainly unsound? Why, that is, are we loathe to accept that the sceptic's conclusion is true? To this the Contextualist has a simple answer: we instinctively know that the conclusion-asserting sentence of [SA] would express a false proposition in a quotidian context in which sceptical hypotheses weren't at issue, and we mistakenly suppose that it's asserting the same false proposition in [SA]. In other words, [SA] strikes us as presenting a profound paradox merely because we're ignorant of what it's really saying, and this because we don't appreciate the indexical nature of knowledge sentences.   (325)

On Schiffer's understanding of it, then, the contextualist's response to [SA] essentially involves

a certain *error theory*—to wit, the claim that people uttering certain knowledge sentences in certain contexts systematically confound the propositions their utterances express with the propositions they would express by uttering those sentences in certain other contexts.   (325)

And Schiffer finds that incredible: 'But that error theory has no plausibility: speakers would know what they were saying if knowledge sentences were indexical in the way the Contextualist requires' (328).

Now, some of Schiffer's construal seems to be based on a misunderstanding of the contextualist's solution. Contextualists need not hold that the skeptic's presentation of [SA] will invariably express the sound argument [SA-T]. I'm the contextualist Schiffer is most directly addressing, and I don't commit to that.[16] What is essential to the contextualist's account, and what I am committed to, is that in her presentation of [SA] (and, in particular, of its first premiss), the skeptic executes a conversational maneuver (and for our current purposes we need not go into the standards-raising rule that the skeptic exploits) that has some tendency to make Tough the operative standards. But whether and under which conditions the skeptic actually succeeds in inducing Tough is a tricky question about which there is room for contextualists to maneuver.[17] To explain why [SA], and in particular, its first premiss, is as plausible as it is, it is sufficient to explain how it is that the skeptic at least threatens to install Tough, at which that premiss is true. The contextualist's ultimate point will

---

[16] Schiffer's reading of me is quite understandable. In section 2 of my (1995), I provisionally assume, for ease of exposition, a 'skeptic-friendly' version of contextualism, and explain how to adapt my explanation for other versions of contextualism, on which the skeptic does not necessarily succeed in raising the standards for 'knowledge' (1995: 6), but execute my subsequent explanations under the provisional assumption. Most of the paper, then, reads as if I accept such a skeptic-friendly version of contextualism.

[17] For discussion of some of the options here, see Chapter 4.

then be that if and to the extent the skeptic's argument does succeed, it does so by at least threatening to install unusually demanding epistemic standards, so its success has no tendency to show that our ordinary claims to 'know' are false.[18] Whether the skeptic actually succeeds in what she threatens may be a very unclear matter.

But none of this helps with Schiffer's complaint; in fact, it may exacerbate the problem Schiffer alleges. What Schiffer finds incredible is that ordinary speakers might be 'bamboozled by their own words': that they might confound what propositions they are expressing by their use of 'know(s)' in dealing with [SA]. It certainly doesn't seem to help matters to add to that that sometimes we contextualist theorists are also quite undecided about what propositions are being expressed!

But there does seem to be an effective response to Schiffer's complaint—disarmingly simple though it is. It is to be found in the fact that if you ask them whether the skeptic's assertion, 'I don't know that I have hands', made as the conclusion of the skeptical argument [SA], contradicts what she would say by asserting 'I know that I have hands' in an ordinary, non-skeptical context, or whether the former denies precisely what the latter asserts,[19] *many ordinary speakers will say 'yes', and many will say 'no'.* (And, interestingly many who say 'no' will go on to volunteer some sort of proto-contextualist analysis of the situation.) Apparently, that's just a tough question. Though some speakers will insist the answer is easy and obvious, this group will split fairly evenly between those who say the obvious answer is 'yes' and those who insist the obvious answer is 'no'. So it simply is the case—whichever is actually the correct answer to the question (if there is a single correct answer: perhaps the answer depends on various details about just how the skeptical argument is presented, whether it meets with resistance, etc.)—that a substantial proportion of ordinary speakers are being 'bamboozled by their own words': they are wrong about whether the relevant skeptical denials of 'knowledge' are incompatible with what ordinary affirmations of 'knowledge' affirm. It's not the contextualist's solution to skepticism that's causing the

---

[18] See again my second 'important point' about contextualist solutions to skepticism at DeRose (1995: 6).

[19] To avoid some needless complications, it's perhaps better to formulate [SA] in the third-person—'Frank does not know that he is not a BIV…', and then ask whether the skeptic's assertion of the third-person conclusion—'Frank does not know that he has hands'—contradicts what would be said by an ordinary, non-skeptical use of 'Frank knows that he has hands'.

trouble. That there is a good deal of 'semantic blindness' going on here is simply a fact that any credible analysis of the situation must face.[20,21] Whether

[20] Stewart Cohen has also answered Schiffer's objection; see his (1999: 77–9; 2001: 89–91; 2004: 191–3). By appealing to how, by means of arguments that are in some ways analogous to the skeptical argument [SA], ordinary speakers can be led, not only to wonder whether certain surfaces they would ordinarily call 'flat' really are flat, but also whether their earlier claims about those surfaces' 'flatness' were true, Cohen persuasively argues that what Schiffer finds so blatantly incredible—that speakers might be so 'bamboozled' about what propositions they are expressing by means of context-sensitive terms—undeniably does happen (assuming, as I take it is safe, that a 'contextualist' resolution to this 'paradox' about flatness is correct). This is a telling point against Schiffer. How can this be so incredible in the case of 'know(s)' when we know such confusion often actually does take place with other context-sensitive terms? Schiffer's professed bafflement here is a dramatic overreaction. But this response of Cohen's is weakened somewhat by a disanalogy between the case of 'know(s)' and 'flat' that Cohen forthrightly admits: 'Having said that, I should note that there is an important difference between contextualist solutions to flatness skepticism and contextualist solutions to justification/knowledge skepticism. Contextualist solutions to flatness ascriptions gain easy and widespread acceptance among most people. But contextualist theories of knowledge/justification do not' (2004: 192; see what immediately follows at 2004: 192–3 for Cohen's response to this problem). Perhaps Schiffer's complaint could be wisely toned down and thereby made a bit more plausible if what he were instead to say is incredible is that there should be a blindness to context-sensitivity when that context-sensitivity isn't quickly and immediately recognized by most speakers once the possibility of context-sensitivity is pointed out to them. Then the disanalogy Cohen admits would be crucial. Thus, it is important in responding to Schiffer to (also) point out that the type of bamboozlement in question happens, not only with *other* terms that are clearly context-sensitive (where this disanalogy applies), *but also in the case at hand (the use of 'know(s)' in skeptical arguments)—whether or not contextualism is right*. Now, at one point (1999: 78), Cohen argues that since contextualism is true, and some speakers resist contextualism, the blindness Schiffer finds incredible really does exist in the case of 'know(s)'. But this presupposes the truth of contextualism, and doesn't point out that the blindness is displayed in the particular situation at hand, where speakers compare denials of knowledge made in the wake of [SA] with ordinary claims to know. That's why it's important to (also) urge against Schiffer the fairly obvious observation I'm pointing out here: That in the case at hand—speakers comparing the content of skepticism-inspired denials of knowledge with ordinary positive knowledge claims—many speakers are wrong about the matter of whether what's being denied in the former is the same proposition that's being affirmed in the latter *whatever the correct account of that relation is*.

[21] In a very interesting passage, Schiffer seems to accept a contextualist account of what we might call 'moderate', non-philosophical skepticism. He begins with an acceptance of some degree of contextualism:

The penumbras of vague terms can dilate or constrict according to conversational purposes. For example, in a conversation about good places to run, a runner might say 'What I especially like about Hyde Park is that, unlike Central Park, it's flat' and count as speaking truly, whereas in a conversation among engineers about where a certain flying device might land in an emergency, an engineer might count as speaking truly when he says of Hyde Park 'It's not flat'. It's clear that knowledge sentences are subject to this sort of vagueness-related variability. In certain conversational contexts you count as knowing that your spouse is faithful; in others you can't really be said to *know*.   (1996: 327)

But, Schiffer explains, this degree of context-sensitivity can't help the contextualist solution to philosophical skepticism; he continues:

The reason this sort of variability is of no use to the Contextualist is that speakers are perfectly aware of when it's going on. If you claim to know where your car is, and someone challenges you, 'But

or not the contextualist solution to skepticism is correct, we simply are stuck with the result that on the matter of whether skeptical denials of knowledge are incompatible with what ordinary knowledge attributions affirm, you can fool a lot of the speakers a lot of the time.

Now, there may be some reason for thinking it's more problematic to suppose that many speakers are blind to the context-sensitivity of their own words than to suppose that many are blind to the context-*in*sensitivity of their own words. But it's not easy to see how to give any credible argument for such an asymmetry, and it's perhaps best not to stretch to anticipate how such an argument might go. And Schiffer does not argue for any such an asymmetry in his paper. In fact, he does not give evidence of even being aware of the fact that many speakers are being bamboozled no matter what the correct answers actually are to our key questions. On the face of it, it looks as though one is involved in equally problematic error theories whether one rules that many speakers mistakenly think there is context-sensitivity when there is none, or whether one rules that many speakers mistakenly think there is no context-sensitivity where it in fact does exist. Either way, many speakers are being 'bamboozled by their own words'.

how can you be sure it wasn't towed or stolen?', you'll merely get impatient at the questioner's obtuseness: it ought to have been mutual knowledge between you that you were speaking casually. The Contextualist who tries to appeal to the context-variability of vagueness must say the following. In a context where scepticism is at issue, the penumbra of 'know' shrinks dramatically to such an extent that one who says 'I know that I have hands' is actually making a false assertion. But the speaker, either because she is unaware of how the vagaries of vagueness affect speech or because she confounds her context, mistakenly thinks that more generous precisifications are in play which count her as speaking truly. This is not a semantic story to be taken seriously.   (1996: 327–8)

But I'm confident that if you ask ordinary speakers, they won't all be so 'perfectly aware' of when this moderate context-shifting is going on in the presenting of such arguments. In fact, if you present them with this 'moderate' analogue of [SA]—

[SA-M] 1-M. You don't know that your car wasn't towed or stolen

2-M. If you don't know that your car wasn't towed or stolen, then you don't know that it's in the Main Street parking lot

So, C-M. You don't know that your car is in the Main Street parking lot

—you will find that fewer, not more, of them accept a contextualist account of the relation between C-M, asserted as the conclusion of this argument, and an ordinary claim of a speaker to know where her car is, than will accept a contextualist account of the original [SA]. In my experience, at least, the biggest difference between the solutions students tend to go for in reaction to [SA] as compared with [SA-M] is that, in the case of [SA-M], more of them go for the 'skeptical solution'—the ordinary claim to 'know' is false (though perhaps appropriate), and the argument shows this to be so—in the case of [SA-M] than in the case of [SA], and thus fewer respondents go for the other solutions, including contextualist solutions, in the case of [SA-M].

Appendix: An Objection to Contextualism from a (Relative) Lack of Clarifying Devices for 'Know(s)'?

In section 5, above, I address the anti-contextualist argument based on how 'know(s)' behaves within speech reports (and based in particular on the DSK principle), that Hawthorne presents at 2004: 98–107. As I remark in n. 5, in one important respect the basic structure of Hawthorne's argument is a bit opaque to me. At 104–7, Hawthorne urges, against contextualism, that we have 'very few' clarifying devices for 'know(s)' (105). Given the context, it seems this isn't intended as an independent argument against contextualism. Having given the meat of his main argument based on DSK, and immediately after writing 'This completes my presentation of the argument' (104), Hawthorne considers this objection: 'But couldn't an analogue of this argument be raised for, say, all context-dependent comparative adjectives? And wouldn't this show that the argument proves too much?' (104). He immediately responds: 'There are disanalogies here that should not be overlooked,' and goes on to explain what he means by 'devices for implementing the clarification technique' and to argue that there are few such devices for 'know(s)'. Thus, it seems that, rather than an independent objection against contextualism, Hawthorne means for these observations about devices of clarification to somehow limit the power of his main argument at 98–103 so that it (or an analogue of it) won't work just as well against the context-sensitivity of some terms that are obviously context-sensitive as it does against contextualism about 'know(s)'.

However, Hawthorne does not explain how to modify the principles he uses in his main argument so that they will show a lack of context-dependence where there is a shortage of clarifying devices but not where there is an abundance of them. I'm supposing we won't want to simply tack 'provided that there are few clarifying devices for "know(s)" in ordinary talk' on to the end of DSK as a weakening proviso. (If Hawthorne's intention is simply to so weaken DSK, then he has to say a lot more about what grounds DSK in the first place: The brief observations on which he does base DSK don't in any obvious way make anything about clarifying devices spring to mind.) So, though Hawthorne seems to present his observations about the lack of clarifying devices for 'know(s)' as a way to *somehow* reduce the scope of his main argument, it is extremely difficult for me to see how it might fill that function in a plausible way.

Perhaps, then, we should consider treating the material at 104–7 as an independent objection to contextualism, based on the (relative) lack of clarifying devices for 'know(s)'. And it's worth addressing such an argument, whether or not Hawthorne intended to be giving it. How powerful would such an objection be?

Well, what is a 'clarifying device'? Though there may be significantly different ways of specifying the details of the notion (and those who wish to use the notion to object to contextualism may try out different precisifications), the basic idea is quite intuitive and straightforward. Since it's also potentially quite important, I will quote Hawthorne at considerable length here. Immediately after writing, 'There are disanalogies here that should not be ignored,' Hawthorne explains:

Suppose I say:

> That is flat.

And suppose that you challenge me by pointing to some little bumps. There are three kinds of tactics available to me.

(i) *Concession.* I concede that my earlier belief was wrong and try to find new common ground: 'I guess you are right and I was wrong. It's not really flat. But let's agree that…'

(ii) *Stick to one's guns.* I claim that the challenge does not undermine what I said. I say (9). You point out some small bumps. I say: 'Well, that doesn't mean it is not flat'.

(iii) *Clarification.* I clarify my earlier claim and then protest that your challenge betrays a misunderstanding of what I believe and what I was claiming. There are various sorts of 'hedge' words that can be invoked in aid of this kind of response.

Here are some examples of clarification:

> *Example 1.* 'The glass is empty'. Challenge: 'Well, it's got some air in it'.
>
> Reply: 'All I was claiming is that it is empty *of vodka*'.
>
> *Example 2.* 'The field is flat'. Challenge: 'Well, it's got a few small holes in it'.
>
> Reply: 'All I was claiming is that it is flat *for a football field*'. (Or: 'All I was claiming is that it is *roughly* flat'.)
>
> *Example 3.* 'He'll come at 3 p.m.'. Challenge: 'He's more likely to come a few seconds earlier or later'.
>
> Reply: 'All I meant is that he'll come at *approximately* 3 p.m.'.

I want to draw attention to the fact that we have very few devices in ordinary life for implementing the clarification technique when it comes to 'knows'.   (104–5)

So clarifying devices for a term would seem to be phrases by which, especially in response to a challenge, one might explain what one meant by the term. (The challenges Hawthorne considers are ones where the truth of a single claim you have made is questioned. Another type of challenge that clarifying devices are helpful in answering is challenges to one's consistency: cases in which you have said one thing using a certain term, and then later, in another context, you appear to say the opposite thing, and your consistency is challenged: 'Hey, but didn't you say yesterday that Reeves *is* tall?' Here we often resort to clarifying devices: 'All I meant then was that he is tall *for a high school player*; what I'm claiming now is that he isn't tall *for an NBA center.*')

And I suppose that if someone thought, like Hawthorne, that there are no or few such clarifying devices for 'know(s)', they might naturally think this is a problem for contextualism, and so they might base an independent objection to contextualism on that observation. The exact nature of the objection could be worked out in different ways, and Hawthorne doesn't help us much here (if he intends an independent objection here at all), but I'm imagining the basic idea would be the thought that it's unlikely that our language would contain a context-sensitive term without also

containing many good clarifying devices by which speakers can indicate the content of their uses of the term on particular occasions. This thought would combine nicely with the observation that there is a lack of such devices for 'know(s)' to yield an anti-contextualist argument.

The problem with such an argument is that the observation on which it's based is wrong, or so it seems to me. At any rate, much more defense would be needed. I think there are many locutions by which speakers clarify what they meant and mean by 'know(s)'. Here are just a few examples (for both first- and third-person uses of knowledge attributions):

'All I was claiming before was that [I know/she knows] it

- quite well
- beyond any reasonable doubt
- by ordinary standards
- by any reasonable standard
- with a high degree of precision.

I never meant to be claiming [to know it/that she knows it]

- for certain
- with absolute certainty
- beyond all possible doubt
- perfectly
- as God would know it.'

For *many* more such examples—and for some very good related discussion—see Ludlow (2005). These seem to be perfectly fine pieces of English, not cases of mere 'philosopher-speak'—though philosophers may well have more occasions to say such things than do other speakers. The speaker's challenger might not like these responses, and might have objections to them, but the speaker doesn't seem to be abusing the English language in using these as devices of clarification.

It's important here that these are not just modifiers that 'S knows that p' can take, but function as devices of clarification: In a context in which one's meaning or consistency are questioned, these devices at least arguably can function to clarify the content of one's claim. Consider for purposes of comparison: While 'quite well', 'with a high degree of precision', 'as God would play it', and 'perfectly' can be used to fill in the blank in 'She played the piano piece _____', the phrases there don't seem to function as devices of clarification: You can't appropriately fend off the kind of challenge we're imagining to an earlier claim to the effect that the relevant person did play the piano piece by means of such 'clarifications'. By contrast, the devices considered above for 'know(s)' do at least seem to function as 'clarifying devices'—at least so far as our current understanding of that phrase goes.

I know: It may be hard to distinguish between an attempt to specify the content of one's use of a genuinely context-sensitive term and an attempt to explain what one meant to say as opposed to what one did say (perhaps while using context-insensitive

terms). Thus, at least on some ways of understanding 'clarifying device', what seems to be such a device might not actually be such. If what we mean by 'clarifying device' is such that a phrase will only count as such a device for a given term if that term is in fact context-sensitive and the phrase is used to specify the content with which the term is being used on particular occasions, then of course it will be very difficult to ascertain that a given phrase is in fact a device of clarification for a given term if it is controversial whether the term in question is context-sensitive. If we instead mean by 'device of clarification' a phrase that *seems* to play the role just described or that *at least arguably* plays such a role, then I claim that there are many such devices of clarification for 'know(s)'. Either way, the fact that there are many locutions that at least seem to play the role in question is an effective counter to the objection to contextualism we are now considering, at least so far as I can see how to work the objection out. It's of course open to potential objectors to try to work out better versions of this kind of objection—versions that, for instance, might contain certain ways of understanding the phrase 'clarifying devices' such that it can be effectively argued both that 'know(s)' suffers a deficiency of such devices and that the deficiency in question is a good reason for doubting that contextualism is correct. All I can say is that I haven't yet seen any version of this objection that seems worrying.

Hawthorne's grounds for denying that there are many such devices for 'know(s)', at any rate, are quite dubious. First he writes:

Granted, we have the locution 'by S's standards'. But in ordinary English this does not work in the manner that the contextualist philosopher intends. An ordinary person might say 'By Hitler's standards, all non-Aryans are bad', but in no way intend to suggest that 'All non-Aryans are bad' was true in Hitler's mouth. (105)

Here, Hawthorne is right that 'by S's standards' doesn't always carry a meaning that would make it suitable for use as a clarifying device. On the other hand, it at least seems that it is *sometimes* properly used as such a device: While it can be used merely to describe what it would take for S to (perhaps quite wrongly) judge that a term applies, it also sometimes seems to function as a clarifying device in a response to a challenge, in which case the phrase seems to mean something quite different.

In any case, 'by S's standards' isn't the best example of a device of clarification for 'know(s)'. Hawthorne then moves on to consider other examples, some better than others:

There are other adjuncts too that are sometimes appended to 'knows': 'He knows, in effect'; 'He knows with reasonable assurance'; 'He knows for certain'; 'He knows full well'; 'He pretty much knows'. But these are rarely put to the service of the clarification technique: If I say 'I don't know' at one time and someone later complains 'You did know', it is not common practice to reply along the following lines: 'It's true that I knew. But what I meant back then was that I didn't know for certain'. (2004: 105 n. 120)

It is true that the attempt at clarification that Hawthorne presents at the very end of the material indented above does not sound terrific (though it's also not awful). But to judge whether there are devices of clarification for 'know(s)', we should give

the proposed devices their best shot at working, and Hawthorne's attempt doesn't do that. With many context-sensitive terms, attempts at clarification go most smoothly when both the earlier and the present content are clarified: directly combining an unclarified object-level claim with a clarifying explanation of one's earlier meaning involving that same term with a different content is often quite clunky (see nn. 12, 14, and 15, above). As I've already said, the examples I give in the text here, which contain two clarifications, do strike me as being quite clear examples of perfectly good English sentences, and, at least so far as I understand what we're meaning by 'clarifying devices', the relevant phrases seem to be functioning as such devices. In the material quoted above, Hawthorne also appeals to how rare it is for such phrases to be used as devices of clarification for 'know(s)'. Hawthorne shouldn't be surprised by that infrequency, because he thinks that situations in which we would have occasion to use such devices—situations in which our consistency involving claims we've made about who 'knows' what in different contexts are challenged—are infrequent: 'Our epistemic practice runs smoothly not because we have clarification techniques available when responding to challenges, but because we are sparing about raising challenges in the first place' (2004: 105). But when such challenges *are* raised, responses by 'double clarification', as we might put it, intuitively seem to be perfectly valid responses—responses that intuitively seem as correct as any of the alternative responses.

I should note that some may develop a more particular worry about contextualism here: that the proposed devices of clarification for 'know(s)' don't seem very precise. It can seem that when the devices of clarification (or at least some of them) are used with, say, gradable adjectives, they quite clearly and precisely indicate what's meant by the adjective on a particular occasion of use. By contrast, it can seem that when one uses one of the proposed devices for explaining what one means by 'know(s)', one is still leaving one's exact meaning quite vague. And one might wonder whether our language would contain context-sensitive terms for which it doesn't also contain devices by which speakers can quite easily *and quite precisely* clarify what they mean by those terms. However, as I argue elsewhere (DeRose 2008: esp. 143–8), this contrast is mistaken: the devices used to clarify uses of gradable adjectives (primarily, 'for an F' locutions, as in 'I meant tall for a fourth-grader') don't really specify the content of the use of the adjective very precisely at all.

I close by again acknowledging that it's of course open to critics of contextualism to try to make good on this kind of argument—to produce a convincing 'objection from a lack of clarifying devices'. I can't anticipate all the ways this might be tried. But as of now, I haven't seen or been able to come up with an objection of this type that seems very worrying.

6

Now You Know It, Now You Don't

Intellectualism, Contextualism, and Subject-Sensitive Invariantism

In these last two chapters, we continue to evaluate objections to contextualism, and we critically compare contextualism with its rival, SSI, concluding that contextualism is the better view. The objections to contextualism we discussed in previous chapters, of course, are relevant to the comparison of the two views, and, indeed, some of those objections have been pressed by proponents of SSI. I am in part relying on my answers to those objections in drawing my conclusion in favor of contextualism here. In these last two chapters, we investigate areas where both contextualism and SSI can at least appear to face problems. In doing so, we will cover what strike me as SSI's two greatest problems. Chapter 7 will concern, among other matters, one of these two problems: SSI's inability to properly handle certain third-person attributions of knowledge. Here in Chapter 6 we are concerned with SSI's other great problem: its denial of intellectualism and its related susceptibility to what we can call 'now you know it, now you don't' objections.

1. Intellectualism, SSI, and Contextualism

As I explained back in section 11 of Chapter 1, along with Stanley, we are using 'intellectualism' to denote the view that the factors in virtue of which a true belief amounts to knowledge are exclusively truth-relevant, in that they affect how likely it is that the belief is true, either from the point of view of the subject or from a more objective vantage point. Thus, to recall the illustration I used back in Chapter 1, the intellectualist can agree with the plausible judgement that, in Ginet's barn example, whether or not Henry knows that he is now seeing a barn can depend on whether he is in a normal situation in

which all the objects that appear to be barns in the vicinity really are barns or is instead in an area teeming with convincing fake barns, such that what he is presently looking at is the only real barn in the region, even if, in the latter scenario, he is completely unaware of the fact that he is surrounded by fake barns. Though this difference has no effect on how likely it is from Henry's own very limited point of view that he is seeing a barn, since he is oblivious to the presence of the fakes in the second case, it does impact the likelihood that Henry's belief is true from a more objective vantage point. Thus, the intellectualist can allow that the difference between our two scenarios can matter to whether or not Henry knows that what he is now seeing is a barn. By contrast, the intellectualist cannot allow that such 'practical' matters as how important it is to Henry that he be right can similarly matter to whether or not Henry knows it is a barn.

We should quickly tighten up the notion of intellectualism a bit, to avoid a potential problem. Intellectualism should be construed so that, in addition to truth-relevant factors, what attitude the subject has toward the proposition in question—how confident she is that the proposition is true, how strongly she believes it—can matter to whether the subject knows. Stanley construes intellectualism such that only truth-relevant factors can matter to whether a subject's *true belief* amounts to knowledge, and if one thinks that the condition that the subject believes that p correctly and completely captures what is required attitude-wise for a subject to know that p, then one will think that, having specified that we are only dealing with true *beliefs* here, we don't have to worry any more about the subject's attitude toward p in our construal of intellectualism. However, it is highly questionable whether belief does so capture the attitude requirement on knowledge, and I, for one, doubt that it does: Perhaps it can happen that a subject believes a proposition, but isn't confident enough in her belief to count as knowing it, and so fails to be knower for purely attitudinal reasons.[1] I am suppressing that

---

[1] *S is certain that p*, I think, is a better candidate than is the ever-popular *S believes that p* for expressing the attitude requirement for *S knows that p*. Such expressions of 'personal certainty' do merely describe the subject's attitude toward p, and do not also evaluate the propriety of the subject's attitude, I believe: No matter how unjustified S is in her belief that p, if she has absolute, unshakable confidence that p is true, you cannot truthfully say, 'S is not certain that p'. (Expressions of '*im*personal certainty'—'*It* is certain that p'—are of course a very different matter. These do contain an important element of epistemic evaluation. For my best attempt to analyze *these* expressions, see DeRose (1998), which is mainly concerned with expressions of epistemic possibility, but at section viii (78–9), explains how to extend my proposed analysis of epistemic possibility to expressions of impersonal certainty by taking 'It is certain that p' to be the dual of 'It is possible that p'.) I am far from certain that even this better candidate perfectly captures what is needed, attitude-wise, to know that p. Perhaps there is no common phrase in natural English that perfectly captures this requirement, in which case we may have to make do with calling it something like 'the attitude of knowledge'.

doubt in this book, mentioning it only here and in a few other places, and otherwise pretending, for ease of exposition, that belief really is the 'attitude of knowledge'.² Intellectualism should not be construed so that it is demanded of all would-be intellectualists that they accept the dubious claim that belief is the 'attitude of knowledge'. So we will instead understand intellectualism to be the thesis that, *in addition to factors concerning what attitude toward the proposition in question the subject has*, only truth-relevant factors can matter to whether the subject knows that proposition.

SSI is an anti-intellectualist view at its very core; it is driven by the thought that non-truth-relevant, practical factors can figure prominently in whether a subject knows something to be the case. What about contextualism?

Contextualism and SSI unite, against classical invariantism, in their rulings about the truth-values of knowledge attributions (and denials of knowledge) in *some* key test cases. In some first-person cases like the Bank Cases presented in Chapter 1, contextualism and SSI agree that the claims to knowledge in LOW are true, and that the admission that the speaker/subject doesn't 'know' in HIGH are also true. How can both assertions be true, given that the speaking subject seems to be in an equally good epistemic position in the two cases? Since the cases are the same with respect to their truth-relevant features, contextualism and SSI both make non-truth-relevant, 'practical' matters relevant to whether the subject *counts as knowing* in her own context, where by this we mean that they make such factors relevant to whether *it can be truthfully said that the subject 'knows'*.³ Contextualism and SSI similarly unite

---

² Perhaps Stanley is similarly pretending when he gives his characterization of intellectualism. For ease of exposition, I myself have sometimes engaged in that pretense in written work in which I didn't state I was doing so.

³ Here I'm stipulating a semi-technical usage of 'counts as knowing'. It is only semi-technical in that this is one thing that the phrase might naturally be taken to mean. In fact, given how I use the phrase in my talks and writings about contextualism, I think it is what most listeners and readers have (correctly) taken me to mean on occasions when I haven't given an explicit explanation of my meaning. However, I am picking out one of several things such a phrase can mean in ordinary English, and when I have successfully used the term to convey this idea even without an explicit explanation, that success has been due to the set-up I provided for my use of the term: I would typically use a few long-winded and very explicitly higher-level statements about the truth-values or truth-conditions of sentences involving 'know(s)', before sliding into the shorter statements that use locutions like 'counts as' to more handily convey the same thoughts. Another thing one could mean by saying that a subject 'counts as knowing' something is, for instance, that they would be (rightly or wrongly) taken to know the thing by a certain group of people or in a certain setting: 'Well, in Iowa I count as knowing that, despite its intractable problems, ethanol is the way to go.' And sometimes, I think, 'this counts as that' is used to express that this *is* that. Thus, for instance, in the quotation we're about to encounter in the next section, when Fantl and McGrath write, 'My true belief cannot count as knowledge, and yours not, simply because you have more at stake than I do in whether p', I think they are expressing the thought (with which they disagree) that a mere difference in stakes cannot yield a situation where one belief does constitute knowledge while another doesn't.

against classical invariantism in the verdicts they issue about the truth-values of assertions involving 'knows' in certain third-person cases—for instance, cases in which subjects are being described as 'knowing' propositions (or not 'knowing' them) in connection with certain evaluations of (potential) actions of the subjects. (Such cases were mentioned back in Chapter 2, and we will focus a lot of attention on them in Chapter 7.)

Because contextualism unites with SSI in this way, it's tempting to think that, like SSI, contextualism is an anti-intellectualist view. But this temptation should be resisted, for it rests on a levels confusion and it overlooks crucial differences in how the two views treat the relevant cases, even though they end up assigning the same truth-values to the assertions made in those particular cases. The contextualist does not hold that whether a subject *knows* or not can depend on non-truth-relevant factors; he holds that whether a speaker *can truthfully describe the subject as 'knowing'*—whether, in our sense, the subject 'counts as knowing' in the speaker's context—can depend on such factors. Whether the speaker can truthfully describe the subject as 'knowing' can depend on such factors, according to the contextualist, because such factors can affect the precise content of the speaker's claim, *not* because they can affect whether the subject is such as to make true the proposition that the speaker is asserting about her. As Stanley astutely notes (2005: 2–3), contextualism results from accommodating the relevant intuitions about the truth-values of the assertions in such cases as these, *while upholding intellectualism*. It is precisely because the contextualist holds fast to the intellectualist assumption that the matter of whether a given proposition ascribing 'knowledge' to a subject is true can't depend on practical, non-truth-relevant matters that he is led to posit that different knowledge relations are denoted in some of the cases in question.[4] For the contextualist, exactly which proposition gets expressed by a knowledge-ascribing sentence will often be affected by 'practical' factors, but the particular proposition that does get expressed will not itself be at all about those factors: Whether that proposition is true is determined only by the subject's attitude and the truth-relevant factors of the subject's situation. On SSI, by contrast, the same proposition gets expressed no matter how the practical facts are arrayed, but the truth-conditions of that proposition are such that whether they are satisfied crucially does depend on the practical, as well as on the truth-relevant, facts of the situation. Roughly, whereas on contextualism a knowledge attribution expresses that a subject (has a true belief

---

[4] At least in part: Contextualists have other reasons for how they proceed here. Most notably, I, for one, am largely guided in my handling of these cases by what happens in other, very different, third-person cases. But an allegiance to intellectualism is certainly part of what has driven me—even when I didn't fully and explicitly realize that such an allegiance was at work.

and) meets such-and-such epistemic standards, where exactly which standards are invoked can be affected by practical factors, on SSI it expresses the thought that the subject (has a true belief and) meets the epistemic standards that are appropriate to her practical situation—and whether *that* thought is true depends in part on what the subject's practical situation is.

2. The Problem with Denying Intellectualism

It is no accident that contextualists tend to uphold intellectualism. Though anti-intellectualist contextualism is a coherent possibility, one of the intuitive attractions of contextualism is that it allows one to uphold intellectualism while delivering certain desired results about key test cases.

And intellectualism is an extremely plausible view, as some of its recent prominent attackers have recognized.[5] Thus, that SSI violates intellectualism while contextualism can uphold it is an important problem that SSI suffers from—and an important relative disadvantage it has in its competition with contextualism. We will presently begin to consider the cost that SSI incurs by its denial of intellectualism. Later, starting in section 7, we will address doubts some might have about whether contextualism is really unscathed by such considerations.

As intellectualism would have it, practical matters, like how important it is to the subject that she be right, seem capable of affecting whether the subject knows only if they have an effect on the subject's attitude toward the proposition in question or on some truth-relevant factor. (In a strange situation in which the fact that the matter is important to the subject constitutes good evidence to the subject that she's wrong about the proposition in question, then of course this importance can be relevant to whether she knows.) The thought that such practical matters might otherwise be relevant to whether the subject knows seems to do considerable violence to the concept of knowledge.

Here are Fantl and McGrath expressing intellectualism (or what they there call 'epistemological purism') in terms very similar to the above paragraph, though they helpfully go on to make the issue vivid by asking us to compare two quite abstractly described subjects who differ only in terms of such practical matters. As I consider this expression of intellectualism, it certainly seems to me exactly right, not just about traditional *accounts of* knowledge, but about knowledge itself: traditional accounts seem right on this score. (Fantl and

---

[5] See especially Fantl and McGrath (2007: 580–1).

McGrath would seem to share this intuitive reaction: They treat intellectualism as something to be given up only under significant pressure. They just think there is significant pressure here, and go on to try to provide it.) But each reader should judge for herself:

According to received tradition in analytic epistemology, whether a true belief qualifies as knowledge depends only on purely epistemic factors—factors that are appropriately 'truth-related.' If my true belief that p qualifies as knowledge while yours does not, this must be because of some difference in our evidence regarding p, the reliability of the processes involved in our beliefs that p, our counterfactual relations to the truth of p, and so on. My true belief cannot count as knowledge, and yours not, simply because you have more at stake than I do in whether p. Raising the stakes may indirectly affect whether one satisfies the belief condition on knowledge (because of one's worrying about the costs of being wrong, for example), but it cannot otherwise make a difference to whether one knows.[6]

Here, in judging the plausibility of intellectualism, we are going by our general sense of what is or can be relevant to whether a subject knows something to be the case, and the thesis of intellectualism seems to do a very good job of articulating what kinds of factors can and cannot matter to that.

3. Stakes and Confidence Levels

The plausibility of intellectualism, which I attempted to bring out in the previous section, depends in part on a certain way of construing what it is for a subject to have a certain level of confidence that a proposition is true. As I will put it, it is aided by a 'stable', as opposed to an 'unstable', construal of confidence.

Recall the Bank Cases from the beginning of Chapter 1. In Case B—the high-stakes, high-standards case—my wife informs me that we are in a high-stakes situation, with much to lose if we assume the bank will be open on the next day (Saturday) but turn out to be wrong, and she raises the possibility that the bank might have changed its hours in the last couple of weeks before asking me, 'Do you know that the bank will be open tomorrow?' Here again is the end of my presentation of that case:

Remaining as confident as I was before that the bank will be open then, still, I reply, 'Well, no, I don't know. I'd better go in and make sure.'

---

[6] (2007: 558). Here Fantl and McGrath seem not to be using 'counts as' in the way I do; see n. 3, above.

Note in particular the stipulation that my confidence is not decreased. Of course, this is a case I'm making up, and I'm free to stipulate features of it. But there's a certain way of thinking of confidence—the more 'unstable' way—that can make my presentation of the case hard to make sense of. One thinking of confidence in this unstable way is likely to object to my description of Bank Case B along the following lines:

> How can it be that he is 'remaining as confident' as he was before?! Before he was aware of the high stakes of his situation, he was happy to just assume that the bank would be open the next day, to flat-out assert that it would be open, and to act on the assumption that it would be open, all without feeling any need to 'make sure' he's right. It seems that, even without further checks, he took himself already to be sure. But once he's aware of the high stakes, all of that seems to change, and we're now told that he is going to do something he's willing to call 'making sure' that the bank will be open before he'll go back to the confident attitude he had before he was aware of the high stakes of his situation. So it certainly seems as if his appreciation of the high stakes of his situation has diminished his confidence quite considerably indeed. So I'm not sure how to make sense of the statement that his confidence is not diminished.

What typically happens to our level of confidence that propositions are true when we find ourselves in high-stakes situations in which we're no longer willing to act forthrightly on those propositions? The above objection seems animated by an 'unstable' conception of confidence (in the truth of propositions) on which it commonly happens that our confidence is very significantly reduced in high-stakes situations. This unstable conception ties confidence closely either to the various types of behavior associated with certainty/uncertainty (Is the subject flat-out asserting the proposition in question? Is she forthrightly acting as if the proposition is true?) or with certain dispositions—which we might call 'local' dispositions—to such behavior (Is the subject presently disposed to flat-out assert the proposition if asked or if it otherwise becomes conversationally relevant? Is the subject presently disposed to act forthrightly as if the proposition is true?).

As opposed to the unstable approach that is described in the above paragraph and that animates the above indented objection, I've always found it more natural to construe confidence, at least in the way relevant to issues of knowledge, in a more 'stable' way according to which our confidence is *not* typically diminished much (if at all) in such high-stakes situations. (Explanation for why 'typically' is included here will be given in the next paragraph.) This more stable construal ties confidence to different dispositions that tend to

remain steady as a subject moves back and forth between high- and low-stakes situations. Whether a subject is in a high- or a low-stakes situation, she will typically have the same dispositions concerning whether she will flat-out assert a proposition should it happen that it is conversationally relevant and she is in a high-stakes situation. So, for instance, even in LOW Bank Case A, I'm disposed *not* to flat-out assert that the bank is open should it happen that that fact is conversationally relevant *and I'm in a very high-stakes situation*. Likewise, when I'm in HIGH Bank Case B, I still have the disposition to indeed flat-out assert that the bank is open on Saturday should it happen that that fact is conversationally relevant *and I'm in a low-stakes conversation* much like that exemplified in Case A. The 'stable' notion of confidence I find plausible here is the one on which the notion more closely tracks these more stable dispositions. On this more stable construal of confidence, if I'm a fairly normal person, I'm about equally confident in the truth of the relevant proposition in LOW and in HIGH, but the same level of confidence that produces confident behavior in LOW doesn't produce similarly confident behavior in HIGH, because that same level of confidence is not confidence enough to yield that behavior, given the gravity of the situation in HIGH. Since I think this more stable construal is the correct construal of confidence (at least as the notion is related to issues concerning knowledge), I'll use scare quotes to refer to 'confidence' as that notion is (mis)construed in the unstable way described in the preceding paragraph.

Though the more stable dispositions to which I'm inclined to tie the notion of confidence *typically* don't change much as a subject moves into a high-stakes situation, it is *possible* for them to be even drastically changed by such a move, and therefore for a subject to be made much less confident, even where confidence is construed in the more stable way, by finding herself in a high-stakes situation. It *could* happen that finding herself in a high-stakes situation really shakes up a subject in such a way that she's no longer disposed to act in the various ways associated with certainty even should she again find herself in more normal circumstances. But this would be a rather unusual occurrence.

The intuitive plausibility that intellectualism enjoys seems to in part depend on, or at least to be aided by, construing the notion of a subject's confidence that a proposition is true in the more stable way in which I'm inclined to construe it—as being such as to track the more stable dispositions, as opposed to what we are calling the more 'local' dispositions. For even the intellectualist can and should hold that non-truth-relevant factors can affect whether a subject knows something to be the case if they do so by affecting the subject's level of confidence that the proposition is true. So the key question dividing intellectualists from non-intellectualists on the relation between stakes

and knowledge is whether two subjects, *both of whom are equally confident* that a proposition is true, can differ from one another in whether they know the proposition in question just because one but only one of them is in a high-stakes situation in which it is much more important that she be right. Now, we can ask this question even if we are construing 'confidence' in the more unstable way in which it tends to track the more local dispositions. But on that construal, when we ask the key question, we are asking what happens in very unusual circumstances: We are asking about subjects, at least one of whose relevant local dispositions is wildly out of whack with what her situation is. And it is wise to be cautious in judging situations that are so far from what is normal. In addition to the doubt occasioned by the strangeness of the situation we are judging, there is also the following reason to think that two subjects of the relevant types might well differ from one another on the matter of whether they know a proposition to be true: It is not only *typical* for our level of 'confidence' (where confidence is construed in the unstable way, as tied to the local dispositions) to vary significantly as we move from a low- to a very high-stakes situation, but it seems that our level of 'confidence' (so construed) *should* vary considerably between the two types of situations. One who didn't become considerably less 'confident' (in this more unstable sense) as she moved into (what she can see is) a very high-stakes situation would not seem to have the appropriate level of 'confidence' toward the proposition. Thus, for instance, one of the subjects we are comparing might be in a low-stakes situation, and might be believing the proposition in question in just the way she ought to. When we then stipulate that the other subject, who is in a very high-stakes situation, is every bit as 'confident' in the truth of the proposition as is the subject in the low-stakes situation, our stipulation forces us to construe the subject in the high-stakes environment as having a very inappropriate attitude toward the proposition in question. And that wildly inappropriate attitude seems capable of resulting in the subject being, or at least seeming to be, a non-knower.

By contrast, when we ask the key question while construing confidence in the more stable way that I think is appropriate to issues of knowledge, we are construing the comparison in a way that involves quite normal subjects, whose 'local' dispositions do vary considerably in reaction to changes in the gravity of their circumstances, but whose dispositions of the more stable variety are not much affected by their being placed in a high-stakes situation. Here, the intuition seems quite strong that these two such subjects—who don't differ from one another in terms of any truth-relevant factor, and who are equally confident in the truth of the proposition—cannot differ from one another on the matter of whether they know the proposition in question.

4. 'Now You Know It, Now You Don't' Problems

To many, the intuitive power of intellectualism becomes most apparent when they consider certain 'now you know it, now you don't' claims, as I've long called them. This class of claims is a bit broader than my label for them would indicate, because they can involve modal comparisons (where what is actually known or not known is compared with what would have been known or not known in a counterfactual situation, or in which two counterfactual situations are compared with one another), as well as temporal comparisons (in which what is known or not known at one time is compared with what is known or not known at another time). They are claims to the effect that matters were (or will be, or would have been) different with regard to whether a subject knows some fact when (or if) such-and-such was (or will be, or had been) the case, where the differences in circumstances being imagined between the two situations being compared (one of which is often the actual, present situation of a subject) concern only some non-truth-relevant matters. Though the same effect is achieved when the difference between the situations concerns other non-truth-relevant issues, our special interest in SSI makes it profitable for us to consider comparisons that allege that knowledge varies with the matter, crucial to SSI, of how high or low the stakes are for the subject; so consider:

> She does know, but she wouldn't have known if more had been at stake
>
> She doesn't know now, but she will know tomorrow, when less will be at stake
>
> She knows on the weekends when she isn't on duty and is only wondering out of idle curiosity; but on weekdays, when much rides on whether she's right, she doesn't know.

Where there is no difference at all in any truth-relevant factors to ground the differences in knowledge alleged in these sentences (where, for instance, there's no difference in the subject's grounds for the belief in question), and there is no difference in the subject's level of confidence in the proposition in question, it is very difficult to believe that such sentences express truths. This is a nasty problem for SSI, according to which the practical matter of how important it is to the subject that she be right is crucially relevant to whether a subject knows, and according to which, it seems, such 'now you know it, now you don't' claims would very often be true about the situations in question (in which there is no difference in the subject's level of confidence in the proposition or in any truth-relevant matters).

This problem for SSI is of course intimately related to the problem we looked at in section 2—so much so that it may be best to view these as two faces of a single problem. In section 2, I was noting and trying to make manifest the plausibility of intellectualism (and implausibility of anti-intellectualism) to our general sense of what is or can be relevant to whether a subject knows something to be the case. But these comparative sentences we are now considering might be able to tempt toward intellectualism those who don't have much of a general sense about what kind of factors can be relevant to knowledge, or who don't trust such a sense. For even those who don't have, or who don't trust, a strong response to a fairly abstract question about what kinds of matters questions of knowledge might turn on might be able to sense that there's something very wrong about these particular 'now you know it, now you don't' sentences. They bring to light the absurdity of thinking that certain types of factors are relevant to the issue of whether a subject knows something. So intellectualism is supported by two closely related considerations: its plausibility to our general sense of what kinds of factors can be relevant to knowledge, and its ability to explain why 'now you know it, now you don't' sentences seem so wrong about the relevant situations.

Briefly returning to the issue I pursued in the previous section, I think considering these 'now you know it, now you don't' sentences also tends to validate construing the notion of confidence, at least as it relates to issues of knowledge, in the more stable way I favor, rather than in the more unstable way in which it tends to track people's more 'local' dispositions to behave in confident/unconfident ways. When I presented the sentences here, I stipulated that they were being asserted about subjects whose level of confidence in the relevant propositions didn't vary between the two situations being compared in the sentences. But when that stipulation is not explicitly made, most people still find the sentences to be very problematic. When we consider these sentences, we imagine the subjects that they concern to have fairly normal psychologies, having not been told that there's anything weird about them. But the relevant 'local' dispositions of normal subjects do vary greatly between low- and very high-stakes situations. So, if we construed the notion of confidence (in the truth of propositions) in the more unstable way in which it tends to track those 'local' dispositions, we would be imagining subjects whose level of 'confidence' in the truth of the relevant proposition varied greatly between the situations being compared. And then, since the matter of how confident a subject is that the relevant proposition is true is clearly relevant to whether she knows that proposition, we probably wouldn't find the sentences to be so problematic. That we find the sentences to be quite problematic even when we're not explicitly instructed to think of the subjects involved as being

equally confident in the two situations, then, supports construing the notion of confidence in the more stable way I favor—in which it tends to track the dispositions that tend to be more stable in normal subjects as they move between low- and very high-stakes situations.

5. KAA to the Rescue?

Hawthorne attempts to handle SSI's problems with 'now you know it, now you don't' sentences much like the ones I mentioned by appeal to the knowledge account of assertion (KAA).[7] As Hawthorne points out, by utilizing KAA, we can explain why sentences like 'I don't know, but I did know before the stakes were so high' are so problematic without invoking intellectualism: If the speaker indeed doesn't now know the proposition in question, as is needed to make the first conjunct of the sentence true, then she also doesn't now know that her past belief in that proposition was true. Thus, since knowledge requires the truth of what's known, the speaker does not now know that she used to know—even if it's in fact true that she used to know. But if she doesn't know that she used to know, then, by KAA, she is in no position to assert that she used to know. On this explanation, the problematic sentence could well be true, but even where it is true, it is unassertable, and asserting it will involve the speaker in falsely representing herself as knowing that she used to know the proposition in question.

The problem with this escape is that it only helps with some of the problematic sentences that need to be handled, leaving other, equally troublesome examples untouched. For one thing, it only helps with first-person examples. But equally problematic are third-person sentences like 'Smith doesn't know, but he did know before the stakes were so high' or 'Smith doesn't know, but he would have known if the stakes hadn't been so high'. According to SSI, such a statement is often true about the relevant situations, where it seems very implausible: where Smith's level of confidence hasn't changed, and there is no change in Smith's evidence for the proposition in question or in any

---

[7] At (2004: 159–60), Hawthorne so addresses a sentence, based on an example Stewart Cohen gives, where one character says to another, 'Okay, Smith knows that the flight stops in Chicago, but we don't' (Cohen himself put a different sentence in his character's mouth). This is a bit different from the sentences I have considered here, because it is comparing different subjects, rather than the same subject at different times or in different situations. Still, this is the same basic problem as I have been considering: The reason the sentence seems so wrong is that the difference between Smith and the characters who are talking to one another is that it is far more important to the latter that they be right about whether the flight stops in Chicago, and there is no difference between them in truth-relevant factors.

truth-relevant matter. And not only is such a statement often true, according to SSI, but, for all we can see so far, it should also be perfectly assertable in many of the situations where it seems so wrong: So long as the speaker herself is in a low-stakes situation where, according to SSI, she does herself know the proposition in question, or, alternatively, if she has such strong evidence that she knows even by the standards relevant to high-stakes situations, it seems that by the lights of SSI, she should be in a position to assert this clunker. And even among first-person examples, this escape only works where it's the speaker's present, actual situation that is the high-stakes scenario in which, according to SSI, she doesn't know. The maneuver seems impotent when we reverse things, imagining that the speaker's present actual situation is the low-stakes scenario about which SSI rules she knows, and it's the alternative (past, future, or counterfactual) situation in which high stakes kill her knowledge, yielding sentences that, about the relevant situations (in which there is no difference in the subject's attitude, or in her evidence or other truth-relevant matters), sound just as implausible as the ones the maneuver can handle: 'I know now, but I didn't know yesterday when the stakes were so high', 'I know now, but I won't know tomorrow, when the stakes will be higher', 'I know, but I wouldn't have known if the stakes had been much higher'. The only way I can hear those sentences as making any sense is if I suppose that, for example, the high stakes in one of the situations being discussed diminish my confidence in that situation. So when I hold my confidence level and truth-relevant factors fixed over the comparisons these sentences concern, the sentences sound extremely implausible.

6. Does SSI Have Good Company in its Misery?

Following suggestions from Hawthorne, Stanley tries to mitigate this problem for SSI by claiming that other theories of knowledge have to endorse 'equally worrisome' sentences of the type in question. The example Stanley gives is that some versions of reliabilism imply that whether a subject knows that an object is a barn can depend on whether and how many fake barns there are in the vicinity, and thus license claims like: 'Henry doesn't know that's a barn, but if there were fewer fake barns around then he would know that's a barn.'[8]

However, I, for one, just don't find that claim so paradoxical. I guess that's because, as I explained back in Chapter 1, it just doesn't seem so implausible

---

[8] Stanley (2005: 113–14). Stanley credits the basic idea to personal communication he received from Hawthorne.

to suppose that whether a subject knows can depend on whether his way of telling is reliable in the particular surroundings he finds himself in, while, by contrast, I do find it hard to believe that whether a subject knows can depend on the practical facts that SSI claims are relevant. I suppose that any who *do* find the statement about Henry and the barns to be as jarringly implausible as are the sentences that cause trouble for SSI will likely avoid versions of reliabilism that would license such sentences about Henry. I think such people would be wrong to find the sentence about Henry so implausible, but, given that reaction, they would be reasonable in taking it to be a strong strike against a theory of knowledge that it endorses such a claim. But, likewise then, SSI's troubles here will be rightly held against the view by most who consider them, and I don't see how this attempt to find company for SSI's misery does much to mitigate that problem.

7. Contextualism and the Advantages of Intellectualism

Thus far, we have focused on the problems SSI encounters in connection with its anti-intellectualism. In section 1, I explained the way contextualism can uphold intellectualism: by making the matter of whether a subject knows some true proposition p, by any given standard for knowledge, depend only on the subject's attitude toward p and truth-relevant features of the subject's situation. But even such 'intellectualist' versions of contextualism (which would seem to include at least most of the forms of contextualism that are actually endorsed) typically make non-truth-relevant factors of the speaker's situation relevant to what proposition the speaker expresses when she says that the subject does or doesn't 'know'. They thus make such non-truth-relevant factors relevant to whether the subject *counts as knowing* p in the speaker's context, in the sense described in section 1: they make such factors relevant to whether the subject can be truthfully said to 'know' some fact by the speaker. Is this any better?

More generally: Does contextualism really enjoy the advantages of intellectualism? Or can those advantages be rightfully claimed only by classical invariantism, according to which neither the matter of whether a subject knows a proposition, nor the matter of whether they count as knowing it (in our current sense), can depend on non-truth-relevant factors? In the remainder of this chapter, I will argue that contextualism does enjoy the advantages of intellectualism.

8. Contextualism and Simple 'Now You Know It, Now You Don't' Sentences: The Apparent Problem and Two Unsatisfying Contextualist Responses

First, contextualism avoids endorsing the problematic 'now you know it, now you don't' sentences that plague SSI. This may not be obvious at first glance, because, though he can insist along with the classical invariantist that non-truth-relevant factors that don't affect the subject's attitude toward the proposition in question can't affect *whether the subject knows* that proposition, the contextualist will admit something the classical invariantist will deny: that, in our semi-technical use of the phrase, such factors can affect whether *the subject counts as knowing* the proposition in a certain speaker's context (where that speaker will be identical with the subject, in the case of first-person knowledge claims). However, as we'll see, despite the feature of the view just admitted, the contextualist does not, and need not, endorse the absurd 'now you know it, now you don't' sentences some apparently think would be licensed by contextualism.

We'll discuss this by reference to an example that has long been the focus of my thoughts about this issue: Palle Yourgrau constructs the following conversation, and rightly observes that 'something is amiss' in it:

Yourgrau's Dialogue

A: Is that a zebra?
B: Yes, it is a zebra.
A: But can you rule out its being merely a cleverly painted mule?
B: No, I can't.
A: So, you admit you didn't know it was a zebra?
B: No, I did know then that it was a zebra. But after your question, I no longer know.
(Yourgrau 1983: 183)

Yourgrau's target is a contextualist version of the relevant alternatives theory of knowledge (RA).[9] However, Yourgrau's complaint would seem to apply equally well to *any* contextualist view according to which A's question manages to raise the standards for knowledge so that B no longer counts as knowing. The general complaint is that B's last line is absurd, but, because contextualism will hold that B no longer counts as knowing the animal is a

---

[9] Specifically, Yourgrau is targeting the contextualist version of RA that Goldman considers at Goldman (1976: 775–8); see especially Yourgrau (1983: 180–1). For more on RA, and especially its relation to contextualism, see sections 13–14 of Chapter 1.

zebra after A raises her question, that last line is what B should or at least could say if contextualism were correct.

But in fact contextualism does not at all license such a line. Before tackling this objection head-on in the following section by explaining the real reason why contextualism would have B say no such thing, in this section we will look at a couple of other, ultimately unsatisfying, ways in which the contextualist might try to wriggle off the hook here. (Most readers would do well to skip ahead to the start of the next section. The rest of this section is being kept in here for the benefit of those readers who just can't get enough of this stuff—and for those who might actually be wondering why I don't respond to this objection in one of the ways I'm about to discuss.)

The first way begins with the observation that there are fairly realistic ways of imagining the circumstances in which Yourgrau's dialogue takes place on which B's last line may well be true. A's bringing up the possibility that the animal in question is merely a painted mule might well shake B's confidence that it is a zebra,[10] or might diminish B's justification for believing it is a zebra. When someone asks you whether you can rule out some alternative to what you are claiming, that is sometimes because they have some special reason to think that the alternative is not all that unlikely. In this case, B might take the fact that A has asked her question about the painted mules to be some reason to think that the animals might well be painted mules. (Perhaps A has heard that, due to a zebra shortage, many zoos have adopted the deceitful practice of using painted mules in their zebra cages?) This could shake B's confidence and/or diminish her justification. And since B's levels of confidence and justification are clearly relevant to whether she knows, anyone—contextualist or not, intellectualist or not—should admit that in the circumstances we are now imagining, B's last line could well be true.

Now, I think it's clear that Yourgrau did not intend for his dialogue to be understood in such a way. Rather, we're to suppose that A has no particular reason, and no more reason than B has, for thinking that the animal is or might well be a painted mule, and that B realizes all this. Yourgrau intended to be raising a problem specific to contextualists, and so would want the circumstances to be such that it's a particular commitment of contextualism that generates the endorsement of the absurd line. We are now in a position to be more explicit than Yourgrau was able to be about the particular commitment of contextualism that is supposed to be generating the problem: that, according

---

[10] Note that this can plausibly diminish the subject's confidence even where confidence is construed in the more stable way I am inclined to construe it, and not just as confidence in the more unstable way that I reject (as these construals are explained in section 3, above).

to contextualism, as opposed to classical invariantism, whether a subject counts as knowing something (in our semi-technical sense) can depend on non-truth-relevant factors of a speaker's context that don't affect the subject's level of confidence. But then it's vital to letting Yourgrau make his intended point that we imagine the case in such a way that A's raising the possibility of painted mules does not diminish B's level of confidence nor constitute evidence that the animal may well be a painted mule. And that's why, being cooperative readers, we understand Yourgrau's dialogue to be such that A has no particular reason to think the animal is or may well be a zebra, and that B realizes this.

But next we note this interesting fact: When we imagine the dialogue transpiring such that A's question *does* greatly diminish B's level of confidence and/or justification (when we imagine it in the way Yourgrau did not intend), that last line *still* sounds absurd. However, we're now imagining that last line occurring under circumstances in which it might well be true. Thus, there is *something* here that makes that last line *sound* absurd even when it is in fact true. And *whatever* accounts for this strange phenomenon of a true claim sounding absurdly false might also be operative when the dialogue transpires in the way Yourgrau intends it to be understood. Or so the contextualist can suggest: Sure, that last line *sounds* absurd, but it sounds that way even where it's just plain true.

And the contextualist might make similar defensive suggestions about variants of Yourgrau's dialogue. Consider, for instance, a third-person variant of the dialogue. In an attempt to sidestep issues about whether A's question might be diminishing B's confidence or justification, the anti-contextualist might utilize a variant of the dialogue in which A and B are talking about some other subject who is not present at their conversation as they say:

Third-Person Dialogue

A: Does Mary know that that's a zebra?
B: Yes, she knows.
A: But can she rule out its being merely a cleverly painted mule?
B: No, she can't.
A: Ah, so you admit now that she didn't know it was a zebra?
B: No, she did know before your question that it was a zebra. But after your question, she no longer knows.[11]

Since we've now placed the subject whose knowledge is in question out of earshot of our dialogue, we don't have to worry that A's raising of her question

---

[11] I utilized just such a third-person variant of Yourgrau's dialogue to sidestep issues concerning whether the asking of the question might be diminishing the subject's confidence at DeRose (2000: 94).

diminishes the confidence and/or justification of the subject's belief. The anti-contextualist will correctly note that B's last line sounds crazy, and will charge that the contextualist will have to endorse it. However, the contextualist could again claim that an endorsement of that last line might not be so bad, since the last line of Yourgrau's original first-person dialogue also sounds absurd even where we're construing the situation so that the last line there is perfectly true, and whatever is making the last line of Yourgrau's dialogue misleadingly sound so absurd even where it is clearly true can also be operating here when we consider the third-person dialogue.

Of course, it's far from clear that whatever it is that accounts for why the last line of Yourgrau's dialogue sounds so bad even where it's true will also account for why it sounds so bad when we imagine the dialogue transpiring as Yourgrau intended it to, or for why the last line of the third-person dialogue sounds so bad. The contextualist is just suggesting a possibility here. Consequently, this is a fairly weak defense. What's more, I believe that the reason the last line of Yourgrau's dialogue sounds absurd even where it is true does not generalize to the other situations,[12] so the escape from trouble for the contextualist that we are now considering does not provide the relief the contextualist needs here.

Above we were considering a way that the contextualist might suggest that it is not a very bad problem if he has to endorse the last line of the dialogues we've been considering. The other attempted escape we will consider in this section is one by which the contextualist argues that he doesn't really have to endorse that suspect line in the first place. The reason the objector thinks the contextualist is committed to the problematic line is that the contextualist will hold that A's question has made the subject in question (B in Yourgrau's original dialogue, Mary in our third-person version) no longer count as knowing that the animal is a zebra. And on the possible escape we're now discussing, the contextualist can accept a version of contextualism on which A's question will not have that effect.

Some contextualists may think that, although the standards for knowledge can vary according to conversational context, it's not all that easy to change the standards, and, in particular, that the mere mention of a bizarre alternative like painted mules is not sufficient to render that alternative relevant. I have

---

[12] That reason is that, even where that last line is true, since B no longer knows that the animal was a zebra (which is needed to make the second sentence of the last line true), B also doesn't know that her earlier claim to know it was a zebra was true (since knowledge requires truth). But then, by KAA, B is not in a position to assert the first sentence of the last line, and in asserting that first sentence, she falsely represents herself as knowing that she used to know that the animal was a zebra. This is similar to explanations we will discuss below in section 10.

a lot of sympathy for the thought that the mere mention of the alternative is not sufficient for making it relevant, but note that in our dialogues we do have more than the mere mention of the alternative. When A mentions the possibility of cleverly painted mules, there are various ways for B to resist allowing that alternative in as relevant. B could reject that alternative by, for example, saying 'Painted mules?! C'mon! That's absurd! Get outta here with that crazy idea!' Following a remark of David Lewis's, a possible contextualist view (and one that we looked at in section 8 of Chapter 4) is that the mentioning of an alternative makes that alternative relevant *if one gets away with making that alternative relevant*, where one fails to get away with making it relevant if one's interlocutor rejects the alternative's relevance in some way like we've just imagined. But in our dialogues, B does not resist allowing the alternative in as relevant. A does seem to get away with making the cleverly painted mule hypothesis relevant. So we have more than mere mentioning in our example; we have mentioning plus some form of non-rejection. (Indeed, beyond mere passive non-rejection, we seem to have some sort of positive acceptance.)

Still, on some possible contextualist views, even that may not be enough. Some contextualists may think that standards can only be raised if elevated standards are somehow appropriate to the context.[13] So, for instance, many contextualists seem to think that higher-than-usual standards are appropriate when the stakes are unusually high, where it's unusually important that the person in question be right. On one such possible contextualist view, mentioning a heretofore irrelevant alternative may make that alternative relevant only if the stakes are high enough to make it appropriate to allow that alternative in as relevant. Some such contextualists may deny that A's mentioning of the painted mule hypothesis, even together with B's apparent acceptance of that possibility, makes that alternative relevant. Such a contextualist can escape Yourgrau's problem by claiming that A has failed to raise the standards for knowledge to such a level as to make B count as a non-knower, and thus B has no business saying she no longer knows.

However, *any* contextualist will allow that under *some* circumstances, the standards for knowledge can be raised to unusually high levels by contextual or practical features of the speaker's situation that are not truth-relevant factors of the subject's situation. And regardless of what a particular contextualist thinks does manage to raise the standards, she will face a Yourgrau-like dialogue, ending with an absurd-sounding last line of the form 'I did know then, but I no longer know now that…', where the line is completed by whatever

---

[13] See for instance the various 'reasonableness' views discussed in section 9 of Chapter 4.

non-truth-relevant factor it is that the contextualist in question thinks does manage to raise the standards. I propose, then, that contextualists not respond in the way we've been considering. So let's just suppose (or pretend, if need be) that A has managed to raise the standards for knowledge in our dialogues to the point at which B and Mary no longer count as knowing that the animal is a zebra. It will turn out that, even so, there really isn't a good objection here.

9. Why Contextualism Does Not Endorse the Simple 'Now You Know It, Now You Don't' Sentences

Contextualism simply does not predict that speakers will say the absurd lines we are considering in the cases in question. The relevant cases are ones in which speakers are initially employing epistemic standards that the subject they are talking about meets, but then move to higher standards that the subject does not meet. (Alternatively, they can begin with the high standards the subject does not meet, but then lower them to a level that the subject does satisfy.) These are indeed cases in which, in our semi-technical sense, the subject first *counts as knowing* in the speakers' context, but then no longer *counts as knowing* in their context (or, in the alternative scenario, first doesn't count as knowing, but then does so count). But that doesn't mean that the speaker should or can say anything along the lines of, 'Well, she knew, but she no longer knows.'

   Charging contextualism with countenancing these absurd lines would seem to be based on the assumption that, in describing someone's past condition, the standards for knowledge are set by what the conversational context was at the time being talked about. But it's open to the contextualist to hold that when a speaker describes a past time, the standards that govern such talk are those set by the speaker's conversational context at the time of the speaker's utterance. Indeed, this option is more than just *open* to the contextualist. Given what happens with at least many context-sensitive terms, we should *expect* the operative epistemic standards to be set by the context at the time of utterance. Though it's possible, I suppose, for there to be a context-sensitive term which operates in the way the objector to contextualism assumes (when describing a past situation, the term takes on the sense that would be determined by the context at the time being talked about), that's not how context-sensitive terms usually seem to work. For instance, in describing my past location, the meaning of my use of 'here' seems at least usually to be fixed by my

present location, not by what my location was at the time being talked about. As I now say in New Haven, 'When I was in Houston last year, David was here', competent speakers of English know that by 'here', I mean New Haven (my present location), not Houston (my location at the time I'm talking about).

'Here' is often grouped with 'I' and 'now' into a special class of context-sensitive terms—the 'core indexicals' as I've heard them called—which in various ways have fairly special properties, and 'know(s)' doesn't belong in that special class. But it's not only 'core indexicals' that behave in the manner in question. Modifying an example of Stanley's,[14] consider, for instance, this dialogue involving 'possible':

A: Is it possible to fly from London to New York City in 30 minutes?
B: No flights available to the public now get from London to New York that quickly. It's not possible to do that.
A: Well, I believe that, though it is not yet available to the public, technology that would allow such fast flights does now exist. So it is possible to make the flight in 30 minutes.
B: Yes, I've heard of that technology, too.
A: Ah, so now you admit that it was possible to fly from London to New York in 30 minutes?
B: No, before you started talking about all existing technology, it wasn't possible to fly from London to New York in 30 minutes, but now it is possible to make the flight that quickly.

In her first line, B is talking about what is possible, given what is available to the public—a perfectly fine use of 'possible'. But then A switches the exact content of 'possible' so that anything that is possible given all existing technology is encompassed—another perfectly fine use. (Depending on what it takes to change the content of a context-sensitive term, A's success here may depend on the fact that B acquiesces in this change in meaning.) In our terminology, flying from London to New York in 30 minutes *does not count as possible* at the beginning of the dialogue, but *does count as possible* after A's switch. Still, while we can probably figure out what she is there *trying* to say, B's last line is intuitively wrong (to put it kindly). Why? Because if, as we are supposing, the content of 'possible' at the end of the dialogue is such as to encompass anything that is possible given all existing technology, whether or not the technology is available to the public, then, since it was in that 'thin' sense possible to fly from London to New York in 30 minutes even before A said

---

[14] Stanley (2005: 53); I present and discuss Stanley's actual dialogue in n. 12 of Chapter 5.

anything, B has no business saying that it 'wasn't possible' to fly that quickly before A mentioned all existing technology. B's use of 'possible', even when she is describing what was 'possible' at past times, is governed by her context at the time of her utterance.

Similarly, the contextualist claims that our uses of 'know(s)'—and its past tense—are governed by the epistemic standards that are in place at the time of a speaker's utterance, not at the time being talked about. Thus, in our dialogues concerning 'know(s)', after A's question raises the epistemic standards of the conversation so that the subject being talked about no longer counts as knowing or as having known the item in question, then B should not say the subject 'did know' before A asked her question. For even at the time before A asked her question, the subject did not meet the higher standards that are operative as B issues her last line.

10. The Fortified Objection: 'What I Said'

'Aha!', it may be (and, I've found, often is) objected: 'While contextualism may not license the last lines of the dialogues we've seen *so far*, the spirit of the objection lives on, for it does countenance the clunking last line of the following':

Fortified First-Person Dialogue

A: Do you know that that's a zebra?
B: Yes, I know it's a zebra.
A: But can you rule out its being merely a cleverly painted mule?
B: No, I can't.
A: Ah, so you admit now that you were wrong when you said you knew?
B: No, what I said—'I know it's a zebra'—was true. But now I wish to add this: I didn't know that it was a zebra.[15]

For isn't it precisely the contextualist's view that B's initial claim to 'know', made when the epistemic standards are relatively low, just is true? So, it seems, the contextualist would have to approve of the first half of B's last line: 'What I said—"I know it's a zebra"—was true.' There, it seems, B just is expressing the contextualist position on the truth-value of her earlier claim. But it would also seem that the contextualist would have to approve of the second half

---

[15] This objection was first raised to me some time in 1990 or 1991 by an anonymous referee for a philosophical journal in an evaluation of DeRose (1992). Essentially the same objection has been raised to me many times since then.

of B's last line as well, for there B is using 'know' (at the object level), and since A has raised the standards for knowledge to a level B does not meet (we are supposing), the contextualist would counsel B to say that she doesn't 'know'—just as she does in the second half of her last line. And, in a similar vein, it can be objected that the contextualist must accept a similarly 'fortified' third-person dialogue, with this last line:

Fortified Third-Person Dialogue—Last Line:

B: No, what I said—'She knows it's a zebra'—was true. But now I wish to add this: She didn't know that it was a zebra.

The contextualist could utilize KAA to account for the wrongness of B's last line in the fortified first-person dialogue: Since B doesn't count as knowing that the animal was a zebra by the high standards that govern her last line (as is needed to account for the second half of that line), she also doesn't count as knowing, by those same high standards, that her earlier (low-standards) claim to 'know' that fact was true, since the truth of p is always required for the truth of 'S knows that p'. Thus, B is in no position to assert the first half of her last line. Indeed, in asserting the first half of her last line, she represents herself as knowing, by the high standards governing her last line, that her earlier knowledge claim was true, which, in turn, commits her to the position that she knows by the high standards that the animal was a zebra, which flatly contradicts what she goes on to say in the second half of her last line. No wonder that line sounds so bad!

However, this appeal to KAA cannot explain what's wrong with the last line of the fortified third-person dialogue, at least on some ways of construing it,[16] nor can it handle an alternative version of the fortified first-person case, reversed so that the speaker was using the higher standards that she didn't and doesn't meet in her earlier claim, but has moved to the lower standards that she did and does meet when she issues this last line:

Reversed Fortified First-Person Dialogue—Last Line:

B: No, what I said—'I don't know it's a zebra'—was true. But now I wish to add this: I knew that it was a zebra.

---

[16] KAA can be used to explain what's wrong with the fortified third-person line if it is somehow being assumed that B is in the same epistemic position as is the subject being talked about with respect to the animals being zebras. But if we assume that, unlike the subject being talked about, B is in a strong enough position to meet even the higher standards for knowledge that are in place as she issues her last line, that last line still sounds awful, and the appeal to KAA we're considering cannot explain why.

These lines sound quite problematic, too,[17] and the appeal to KAA we have been considering won't explain what's wrong with them.[18]

[17] There does seem, at least to me, to be a perceivable difference in the exact kind of problem involved here: The last line of the original fortified first-person dialogue seems to quite sharply 'clash', in the way discussed in Chapter 3: It sounds like a contradiction. And this contrasts with the last line of the reversed first-person dialogue, and of the fortified third-person dialogue, which do not in the same way seem to clash. And I believe this difference in appearance is one rightly accounted for by the KAA explanation we've been considering, according to which, in the case of the original, but not the reversed, version of the line there is a contradiction between something B asserts (namely, that she didn't know that the animal was a zebra (by the standards that govern her last line)) and something that B merely represents as being the case (namely, that she does know that same fact by those same standards).

I realize that in cases like these—where the sentences involved are a bit complicated, as opposed to more simple clashes like 'Dogs bark, but I don't know that they do'—many will not be able to discern clashes from other problems, and will not find this perceivable difference in how the original and the reversed version of the last line of the fortified first-person dialogue sound to them: To many, both of these lines will just sound bad, with neither of them any more than the other seeming to 'clash'. And some of these folks may be suspicious of my sorcerous ways when I claim to detect a difference here, suspecting that I don't have a sense for this difference that my theory (KAA) then accounts for, either. Rather, the suspicion will go, it's because I'm theoretically committed to there being a clash here that I tend to hear the one, but not the other line, in a 'clashing' way.

Indeed, in reporting this difference in how these two lines sound, I do feel a bit like one of the few remaining practitioners of an ancient, dying religion. But in a couple of seminars I took from Rogers Albritton at UCLA, I was initiated into the (or at least a) practice of discerning clashes from other ways that conjunctions can be problematic. So—to use examples that we discussed in the seminar, inspired by White (1975), which someone in the seminar was reading, and that I have since dealt with (DeRose 1991: 601−5)—there is to me a discernible difference between

(a) It is possible that X Vs, but it is not possible for X to V

and

(b) It is possible for X to V, but it is not possible that X Vs.

To use a couple of advanced terms of art, while (a) clashes, (b) is just clunky. I have an explanation (based on KAA) for why (a) would clash (see DeRose 1991: 604), but the difference between these two is something I can just perceive (hear)—or so I think. That (a) clashes, and does not just clunk like (b) does, functions for me as a datum to be explained. (Similarly, though David Sosa's '*P* but I don't know whether I know that *P* and the conjunctions Sosa considers later in (forthcoming) seem odd, they do not clash). How did I come to have this ability to perceive such differences? After hearing Albritton and some more advanced graduate students discuss many examples for less than half an hour, I felt like I could join in: 'Yes, this one clashes; that one is just a bit clunky; nothing at all wrong with this other one.' And I don't think I was following any theories in making the distinctions involved: Those that I ruled were clashers really did just come to strike me in a different way from the others. Of course, one can't help but start formulating theories, or at least proto-theories for why some clash and others don't. But I do not think my reactions were being guided by the proto-theories I was inclined toward, because there were surprises: cases where one gets a result that goes against the prediction one's proto-theory would issue: 'Oh, that one does clash. That's interesting. I wouldn't have expected that.' Afterward, I went back to the examples in my notes, and found that the distinction I was inclined to draw between those that clashed and those that didn't aligned fairly closely with other differences between the two camps that were formed. For instance I could come up with situations about which (b), except for being an awkward way of putting things, seemed a good description: in this case, situations in which it is clear that X has the capacity to V, but it is certain to the speaker that X does not exercise that capacity. By contrast, I couldn't come up with situations about which (a), though it is logically coherent, seems a good description. There indeed can be cases in which X in fact does not have the ability to V, but in which it is uncertain to the speaker whether X V's (presumably in part because it is also uncertain to the

To explain with any confidence exactly how these last lines go wrong would require me to be far more confident than I am about the semantics of sentences involving such phrases as 'what I said' (which seem quite similar to the likes of 'I never said that', which we looked at back in Chapter 5), and, especially, how such locutions work where context-sensitive terms are involved in what was said at the earlier time. Nevertheless, we can see that the wrongness of these last lines doesn't provide a good reason to doubt contextualism about 'know(s)' by observing that analogous lines involving obviously context-sensitive terms sound wrong, as well. Consider these last lines (the reader should by now be able to construct the rest of the dialogues that lead up to these last lines for herself):

No, what I said—'It's not possible to fly from London to New York in under 30 minutes'—was true. But now I wish to add this: It was possible to fly from London to New York in under 30 minutes.

No, what I said—'Reeves is tall'—was true. But now I wish to add this: Reeves wasn't tall.

None of these fortified last lines—the two immediately above and the ones we've been discussing involving 'know(s)'—are things any competent speaker would say. Though the first half of such a statement seems to be a claim about the truth-value of an earlier claim involving the relevant key term, and only the second half involves an object-level use of that term, with the possible

---

speaker whether X has the ability to V). However, it seems somehow wrong, and not just awkward, for the speaker to assert (a) even in such scenarios.

Though the last lines of the third-person fortified dialogue and the reversed first-person fortified dialogue don't clash, they do sound (in some other way) quite bad. The defense of contextualism I'm about to give should be sufficient for those who hear no difference between these lines and the last line of the original fortified first-person dialogue. But those who do possess 'golden ears' and discern a special problem with the last line of the original fortified first-person dialogue can account for this clash by the explanation, involving KAA, that I've already given.

[18] This objection is originally due to Shelly Kagan. In a series of talks I gave at various philosophy departments in the years 1996–2001, I presented the objection based on the fortified first-person dialogue, together with my response based on KAA, thinking, at least when I first gave the talk, that that solved the problem. However, at one of these talks, at Yale University, Shelly Kagan raised what I'm here calling the Reversed Fortified First-Person Dialogue, plausibly objecting that its last line also sounds very implausible, and that the KAA-based account cannot handle it. I *believe*, though I'm not certain of this recollection, that Kagan claimed that the last line sounded *just as bad* in its reversed form. I do (and did, as I recall) disagree with that, for the reason given in the previous note. However, it does sound very implausible, and the fact that my KAA-based solution to the original fortified objection couldn't explain what's wrong in the reversed dialogue cast a considerable shadow over that account. Thus, when much of the material of the talks eventually found its way into print in my 2002a (which constitutes the basis for Chapter 3 of this volume), the material on the fortified dialogue was cut out and shelved until I could come to grips with the last line of the Reversed Dialogue.

exception of some examples involving core indexicals, it still seems *somehow* wrong to so jam together, without clarification, two such statements where there is a crucial difference in content between what the key term meant in the earlier statement and what its current content is. Instead of the above, speakers should say something very roughly along the lines of the following (though there are many different ways of formulating such attempts at reconciling earlier claims with what is to be said relative to different standards that are in play as a statement is made):

> No, what I was claiming before was that nothing currently available to the public would allow one to fly from London to New York in under 30 minutes, and that was and still is true. But if we're now talking about what's possible relative to all existing technology, whether or not it's available to the public, then, yes, it is in that sense possible to fly from London to New York in under 30 minutes.

> No, what I was claiming before was that Reeves is tall for a high school player, and that was and still is true. But he certainly isn't tall for an NBA center.

Even where clearly context-sensitive terms are involved, the analogues of our problematic last lines, which contain no such clarifying explanations of the differences in content, seem quite wrong to assert—even though it seems they might well be true. And I would add that for me, and for many others, the appearance of wrongness feels at least roughly comparable to the appearance of wrongness in the last line of our fortified dialogue concerning 'know(s)'. (But reactions here vary in ways I will discuss two paragraphs below.) Thus, it should be no surprise, even if 'know(s)' is context-sensitive, that we find it behaving similarly, so that the last lines of our various fortified dialogues involving 'know(s)' also sound wrong.

Well, it may be asked: What could B say instead in the case of 'know(s)'? How could she use 'clarifying devices' to *correctly* express something along the lines of the somehow incorrect last lines of our fortified dialogues? I think there are many locutions by which speakers clarify what they meant and mean by 'know(s)', and many of them would work just fine here. For several examples and discussion (and also a reference to Ludlow 2005, where many more examples can be found), see the appendix to Chapter 5. But the objection to contextualism we're now considering is that the last lines of the fortified dialogues—lines that don't employ any such clarifying devices—sound very wrong in a way that we wouldn't expect given contextualism. In response, we are noting that analogous lines involving terms that are clearly context-sensitive sound quite awful as well. These

objectors need to revise their expectations about the behavior of context-sensitive terms.

I should convey that some (but far from all) who have reported their reactions to fortified last lines involving 'know(s)' and their analogues involving less controversially context-sensitive terms seem to find varying degrees of wrongness here. Perhaps over-generalizing a bit, the pattern seems to be this: the more obvious is the context-sensitivity of the key term involved, the more likely it is that *some* respondents will find the last line to be less egregiously wrong than other examples. Indeed, in the extreme case of 'core' indexicals, many (but far from all) respondents don't find the relevant lines to be wrong at all.[19] Part of the problem, of course, is an issue concerning whether listeners are understanding the material that is inside the dashes to be a higher-level metalinguistic statement. Much as I dislike the gesture, I tried the examples both with and without making 'air-quotes' with my fingers at the relevant time. I tried it both with and without writing the sentences on the blackboard and then also saying them. I tried it both with and without inserting the word 'namely' right before the quoted material. Still trying to give the objection every chance, I also tried the lines while completely omitting the material inside the dashes—making do with just 'what I said', trusting the respondent to remember what was said earlier in the dialogue. None of this seems to yield what our objector needs. With the exceptions of some comparisons involving core indexicals, there seems to be no good way of constructing the last lines so that they are even close to seeming clearly wrong where 'know(s)' is involved, but clearly all right in the case where another context-sensitive term is used. That would appear to be because these are all just stupid ways of speaking. Frankly, many respondents just professed to be confused: 'Huh?!'

At any rate, I don't think there's any objection worth worrying much about here. Competent speakers just won't say anything like these last lines. When we imagine speakers incompetently saying such things, we *at least* very often find the resulting assertions quite problematic—*even when the assertions involve genuinely context-sensitive terms*. To the extent that the term involved is obviously context-sensitive, we can perhaps sometimes see past the problematic nature of the assertion to judge that it is nonetheless true. But the sense that there

---

[19] It seems to promote the response that the lines aren't wrong at all if the examples involving the core indexicals are presented first. If they're presented after analogous lines involving other terms, respondents are more inclined to find the last lines involving core indexicals to be at least to some degree wrong. I should also note that the last lines involving core indexicals seem to go down more smoothly if unusually heavy stress or intonation is put on the indexical term: 'What George said—"I'm tired"—was true. But now I wish to add this: *I* wasn't tired'; 'What I said—"The book is here"—was true. But now I wish to add this: The book wasn't **here**.'

is something wrong with an assertion of the type in question that involves a term that is not so obviously context-sensitive should not tempt us to conclude that the term isn't really context-sensitive, if a similar intuitive appearance is produced by the assertions that involve genuinely context-sensitive terms. What competent speakers do is utilize clarifying explanations about what they meant and mean by the terms in question on different occasions. And when effective clarifying explanations are used in the examples involving 'know(s)', the resulting claims don't give any strong appearance of being wrong. On the contrary, though invariantists might have theoretically motivated objections to the resulting claims, they *seem* to be perfectly fine examples of good English sentences that explain at least to some extent what a speaker takes herself to have meant and to be presently meaning by her claims involving the term in question. (Again, see the appendix to Chapter 5 for a discussion of clarifying devices for 'know(s)'.)

11. Elusive Knowledge?

As we saw in section 9, contextualism won't counsel speakers to say things like 'First I knew, then I didn't' in the relevant situations where those lines sound absurd—situations in which there's been no change in truth-relevant factors between the two times being talked about. Since there seems to be considerable confusion about the matter, it's worth making the closely related point that contextualism doesn't imply that any knowledge appears or vanishes as conversational context changes.

It seems to many that contextualism renders knowledge unstable or elusive in the sense that it would make our knowledge come and go—be gained or lost—as conversational context changes. And to many, this can seem a very problematic implication of contextualism, since it seems to them that knowledge in fact doesn't appear and disappear due to changes in conversational context. I agree: Knowledge isn't in that way elusive. But, I insist, contextualism doesn't make knowledge elusive in that way. Here I have to contend with friends as well as foes of contextualism. One of contextualism's most prominent advocates, David Lewis, in his contextualist manifesto 'Elusive Knowledge' (Lewis 1996), makes many comments about the elusiveness of knowledge, including the following (emphasis added in all cases), which I've numbered for ease of reference. One bit of background: Lewis contends that engaging in the practice of epistemology typically has the effect of raising epistemic standards. He writes:

1. Maybe epistemology is the culprit. Maybe the extraordinary pastime *robs us of our knowledge.* Maybe we do know a lot in daily life; but maybe when we look hard at our knowledge, *it goes away.*   (1996: 550)

2. In the strict context of epistemology we know nothing, yet in laxer contexts we know a lot.   (1996: 551)

3. Unless this investigation of ours was an altogether atypical sample of epistemology, it will be inevitable that *epistemology must destroy knowledge.* That is how knowledge is elusive. Examine it, and straightway *it vanishes.*   (1996: 560)

4. Imagine two epistemologists on a bushwalk. As they walk, they talk. They mention all manner of far-fetched possibilities of error. By attending to these normally ignored possibilities *they destroy the knowledge they normally possess.* (1996: 565)

I certainly wouldn't want to say any of the above. (Lewis, though he does write all the above, does in a way seem to take it back at the end of his essay, as we'll see below in the following section.) If contextualism really implied that knowledge vanishes (3), is destroyed (3, 4), goes away (1), or is robbed from us (1) by means of conversational developments (often in conversations we're no party to) that have no effect on how strongly we're positioned with respect to the beliefs in question, then I'd be very tempted to reject contextualism. Thus, I can't be surprised that others, who think contextualism has such implications, reject the view on those grounds. And I can't be surprised that foes of contextualism think the view underwrites remarks such as those above, when one of contextualism's strongest friends is the one making the remarks.

But let's get clear, once and for all, about contextualism's implications as the standards change so that a subject who used to meet them no longer does. To increase clarity, let's consider a case involving third-person, rather than first-person, knowledge-ascribing sentences. So: Fred's over there, very confidently believing the true proposition p. Far away from him, and quite out of earshot, Wilma and Betty are talking about whether he 'knows' that p. Contextualism is true (we'll suppose). To make things uncomplicated, let's suppose that just two sets of standards come into play in Wilma and Betty's conversation—'low' and 'high'—and that throughout the story the strength of Fred's epistemic position with respect to p remains constant at a level at which he meets the low, but not the high, standards. Initially, Wilma and Betty were in a context that selected the low standards, so they spoke the truth when they said, 'Fred knows that p'. Now, however, their conversational context has changed; they're now in a high-standards context. So Wilma and Betty can no longer truthfully say that Fred 'knows', and in fact can now truthfully

deny that Fred 'knows'. Let's suppose they have now done just that: They have truthfully said, 'Fred doesn't know'.

What has Fred lost? What has gone away, been destroyed, been robbed from him? 'High' knowledge? No, he never had that. 'Low' knowledge? No, that he still has.

Knowledge? Knowledge simpliciter? Well, if what's meant by 'knowledge simpliciter' is knowledge according to the one-and-only standards that ever govern attributions or denials of knowledge, then, of course, there's no such thing on the assumption that contextualism is correct. (And if there were such a thing—if contextualism were wrong—then *wherever* those one-and-only standards are set, whether those uniquely correct standards are high or low, Fred has not lost this simple knowledge by means of Wilma and Betty's conversational shenanigans: he either has it all along, or never had it.)

If you mean knowledge by *our* current standards, then, of course, it's hard to say whether Fred has knowledge or lacks it, in part because it is difficult to discern what our current standards are in this rather funny context in which I am discussing contextualism in writing, and in part because my example has not been very fully described. All I've specified so far about Fred is that he met (and meets) the 'low' standards set by Wilma and Betty's old context, that he fails to meet their current 'high' standards, and that the strength of Fred's epistemic position didn't change: he went from meeting to failing to meet the speakers' standards due entirely to a change in their standards, and not at all because of a change in the strength of Fred's epistemic position. No matter. We needn't bother with more detail, nor need we fret over what our current standards are. Though we can't yet tell whether or not Fred knows (whether or not he meets our current standards), we are already in a position to say whether or not Fred has lost knowledge of p. He hasn't. Whatever our standards are right now, and however strongly or weakly Fred is positioned with respect to p, given only that the strength of his epistemic position hasn't changed (and *that* we have been given), we can ascertain that Fred hasn't lost knowledge. He of course hasn't changed from meeting the standards we are currently employing to failing to meet those same standards because of anything that transpired in Wilma and Betty's conversation! He either knew and knows, or didn't and doesn't. Either way, he hasn't lost his knowledge that p and no knowledge has gone away, been destroyed, or been robbed from him.

What Fred *has* lost is a certain relation to Wilma and Betty's context. He has gone from meeting the standards they were utilizing to failing to meet the standards governing their conversation. Using phrases like 'Wilma and Betty's standards' fluidly, so that we may speak of those standards as changing

over time (rather than as rigidly referring to the particular standards they are employing at a particular time), he has lost the property of satisfying Wilma and Betty's standards. But because this 'loss' is due to a change in Wilma and Betty's standards, and not at all because Fred's epistemic position has been diminished, this 'loss' doesn't mean that there is any variety of knowledge that Fred has lost.

12. Lewis and Semantic Ascent

At the close of his paper, Lewis acknowledges that there's something fishy about his saying such things as that knowledge vanishes or that we first know and then fail to know:

But wait. Don't you smell a rat? Haven't I, by my own lights, been saying what cannot be said? (Or whistled either.) If the story I told was true, how have I managed to tell it?…I said that when we do epistemology, and we attend to the proper ignoring of possibilities, we make knowledge vanish. First we do know, then we do not. But I had been doing epistemology when I said that. The uneliminated possibilities of error were *not* being ignored—not just then. So by what right did I say even that we used to know?

Lewis then admits that he 'bent the rules' in making those comments—and I presume the same would go for the other comments we've been looking at. He then goes on to explain:

I could have had my say fair and square, bending no rules. It would have been tiresome, but it could have been done. The secret would have been to resort to 'semantic ascent'. I could have taken great care to distinguish between (1) the language I use when I talk about knowledge, or whatever, and (2) the second language that I use to talk about the semantic and pragmatic workings of the first language. If you want to hear my story told that way, you probably know enough to do the job for yourself. (1996: 567)

And we do have to do the job for ourselves, for 'Elusive Knowledge' comes to a close just one sentence later.

But it's not hard to engage in a little semantic ascent and find substitutes for the object–level phrases I've been objecting to. Fred's knowledge doesn't *vanish*. He doesn't first know and then fail to know. Rather, to be quite tiresome, we can say:

> (A) First, Fred was such that the proposition expressed about him by the sentence 'Fred knows' in Wilma and Betty's conversation was true of him. But then, because their context changed so that 'Fred knows' came

to express a more demanding proposition, Fred was such that the (new) proposition that would have been expressed by 'Fred knows' in Wilma and Betty's context was not true of him.

But we don't have to be that tiresome. There are devices for engaging semantic ascent that don't involve explicitly putting knowledge-attributing sentences in quotation marks and explicitly saying things about the truth-values or truth-conditions of those sentences. With just a little set-up, I think one can, as I've been doing in this very book, achieve that same effect by simply saying things along the lines of:

(B) First Fred met the epistemic standards set by Wilma and Betty's context and then, because their standards went up, he failed to meet the standards set by their context.

Or, to be even less tiresome, we can say:

(C) First Fred counted as knowing, then he didn't,

which really is only minimally more tiresome than is 'First Fred knew, then he didn't'. In this book, I've tried to be explicit about my semi-technical use of 'counts as knowing', but in papers and talks, I have used the phrase in that sense without explicit explanation, and it seems that, by first saying something more explicit about satisfying the truth-conditions of particular uses of sentences involving 'know(s)', and then switching over to talk about what 'counts as knowledge', this use has been successful: Listeners and readers have understood me correctly. And, to explicitly deal with the numbered passage from Lewis's 'Elusive Knowledge' that I haven't mentioned since displaying it, instead of Lewis's

2. In the strict context of epistemology we know nothing, yet in laxer contexts we know a lot,

we could instead say the only slightly more tiresome,

(D) In the strict context of epistemology we count as knowing nothing, yet in laxer contexts we count as knowing a lot.

Lewis explains that he is engaging in an 'expository shortcut, to be taken with a pinch of salt' (1996: 566). But contextualists can utilize the likes of (B)–(D) in expository shortcuts to achieve the same effect as saying such truly tiresome things as (A). (B)–(D) build the tiresome 'semantic ascent' of the likes of (A) into easy little phrases like 'standards for knowledge', and 'counts as knowing', which, when spelled out, amount to higher-level claims

about the truth-conditions and truth-values of the relevant sentences. I've found that, in presenting contextualism, one or two (A)-ish statements (in which knowledge-ascribing sentences are explicitly mentioned and something about their truth-conditions is explicitly stated) at the start of a paper or talk is all that's needed for listeners or readers to be ready to correctly understand the likes of (B)−(D). That way we can quite painlessly avoid rule-bending (breaking?) shortcuts of the likes embodied in phrases like Lewis's 1−4.

13. The Fallacy of Semantic Descent

But why shouldn't we use rule-bending shortcuts like Lewis's? When I sent a copy of the then just published DeRose (2000)—the paper on which much of this chapter is based—to him, Lewis wrote back to me, defending his use of such locutions in quite strong terms:

Competent readers know (except perhaps when they're in a philosophy seminar) that they're expected to interpret what's said in such a way as to make the message make sense. So when they read (outside the seminar) that genes are selfish, they know to take it as a metaphor. And when they read (outside the seminar) that France is hexagonal, they know that what's meant is not that France is a perfect hexagon. And when they read (outside the seminar) that tomorrow never comes, because when it comes it will no longer be tomorrow, again they understand perfectly well. And so on. Authors are entitled to rely on readers to exercise just this kind of interpretive competence. But all too many philosophers seem to abandon their interpretive competence when they enter the philosophy room. They turn themselves into literalists and simpletons—or, more likely, they pretend to be literalists and simpletons when really they know better. I think that if this silly habit leads them to misunderstand—or more likely, pretend to misunderstand—that's their problem. Instead of pandering to such silliness, and reproaching those who refuse to do likewise, you might better reproach the silliness itself.[20]

In a more rational philosophical world, there might be no harm in using locutions like Lewis's. But in our less reasonable actual world, the use of such shortcuts seem to illicitly cloud the evaluations of contextualism of many philosophers in a way I'll now try to explain.

---

[20] Personal correspondence of 9 November 2000, copyright The Estate of David K. Lewis, used by permission of Stephanie R. Lewis. I very much hope that much of Lewis's philosophical correspondence will be published in the near future. I benefited greatly from my exchanges with Lewis, and am anxious to see what he wrote to others, especially since I understand he had many philosophical correspondents, several of whom he wrote to quite extensively.

We might call the mistake of holding it against a theory or view that it has a certain implication, when in fact its real implication is not what one thinks, but is rather some related higher-level statement, the 'fallacy of semantic descent'.[21] Theories according to which various types of sentences are held to be context-dependent in meaning are often subject to objections that depend on this fallacy. To get a feel for the fallacy in action, consider a fanciful case in which the higher-level statement (the actual commitment) is uncontroversial. I assume we all believe that the phrase 'this room' works in such a way that the following is true:

(E) It can happen that one person can truthfully say, of a physical object, A, 'A is (wholly) in this room,' while, at the very same time, another person (who is in another room) can truthfully say of the very same physical object, 'A is not (wholly) in this room.'

Now, suppose that someone objected to this view of yours by first claiming that, according to your view,

(F) One and the same physical object can at one and the same time both be and not be wholly in this room,

and then going on to suggest that (F) is absurd, counter-intuitive, or implausible. The proper response to this objection, I take it, is to reject (F), to point out that it's (E), and not (F), that you accept, and to point out that (F) does not follow from (E). To hold it against you that you are committed to (F) is to commit the fallacy of semantic descent.

It is not always easy to determine whether the fallacy of semantic descent is being committed: Potential candidates seem to positively revel in illicitly lurking in the shady borderland between holding the apparent object-level 'commitment' against the theory and just using the object-level statement as a rule-bending shortcut to refer to the higher-level statement. That is, it is sometimes unclear whether the objector is really confusing his target's commitments or whether she is rather, to use Lewis's words, using the object-level statement as an 'expository shortcut, to be taken with a pinch of salt' to express her target's commitments. In our above example, we might wonder whether our objector really thinks we believe or are committed to (F), or whether she's just using (F) as a rule-bending device by which she hopes to express our real commitment, (E). After all, we could—though it seems inadvisable—more or less explicitly *agree* to use (F) as a handy piece

---

[21] Unger calls this fallacy, or at least one much like it, the 'fallacy of conflating contexts' (Unger 1995: 16).

of rule-bending shorthand by which we'll express (E). But in that case, we should all remember that that's what's going on, and we should be careful not to come to doubt (E) because of how bad (F) sounds or how absurd (F) seems. After all, we had to break the rules to use (F) as a device for expressing (E). If (E) implied something that seems absurd, that would be grounds for doubting (E). But it casts no legitimate shadow on (E) that an absurd-sounding statement can, through a rule-bending piece of shorthand, be used to express (E). At any rate, in the absence of any such explicit agreement, we might be left to wonder whether our objector is really committing the fallacy of semantic descent, or whether he rather just considers (E) itself to be crazy. The problem with the first interpretation is that one is ascribing blatantly bad reasoning to the objector; the problem with the second interpretations is that it's hard to see why (E) itself would seem crazy to someone—unless it was being somehow confounded with F.

Moving to cases about which someone might actually be misled by our fallacy, let's now consider a sample published example, whose target is not contextualism regarding knowledge-ascribing sentences, but is rather a theory according to which a certain type of modal expression is held to be such that its 'meaning is relative to the person who uses it'. In his essay 'Certainty', G. E. Moore (the target) considers statements of the form 'It is possible that p is true'. Engaging in some not-all-that-tiresome semantic ascent (note his talk about the meaning, truth-conditions, and truth-values of such statements), Moore writes the following:

The expression 'It is possible that p is true' is, though it looks as if it were impersonal, really an expression whose meaning is relative to the person who uses it. If *I* say it, that I should not know that p is false, is a necessary, though not a sufficient, condition for the truth of my assertion; and hence if two people say it at the same time about the same proposition it is perfectly possible that what the one asserts should be true, and what the other asserts false: since, if one of the two knows that p is false, his assertion will necessarily be false; whereas if the other does not know that p is false, his assertion may be, though it will not necessarily be, true. (Moore 1959: 241)

Now Moore's main claims in the above passage are pretty clearly correct.[22] But Alan R. White objects to Moore's account, writing, for instance, of 'Moore's mistake…of supposing that, since what is not known to me may be known to you, the same thing may be both possible and impossible' (White 1975: 52). But Moore himself would not say or write that 'the same thing may be

---

[22] For a vindication of Moore's claims, see DeRose (1990: 58–72).

both possible and impossible'. At least I've never caught such a statement in his writing. What Moore does write is:

> (G) If two people say it ['It is possible that p is true'] at the same time about the same proposition it is perfectly possible that what the one asserts should be true, and what the other asserts false.

And there are cases that establish that he's absolutely right about that.[23] And (G) simply doesn't imply any statement like

> (H) The same thing may be both possible and impossible

I'm uncertain whether White is committing the fallacy of semantic descent here. I'm inclined to think that he is—largely because I'm unable to come up with a reading of his reasons for rejecting Moore's views here which makes them any more convincing. (But read the first fifty-two pages of White 1975 and judge for yourself.) That is, I suspect that White thinks that (H) is something Moore is committed to, and that he denies Moore's theory in part because (H) sounds implausible to him. But maybe I'm wrong and White realizes perfectly well that Moore is not really committed to (H). Perhaps White is just 'bending the rules' a bit and using (H) as an 'expository shortcut' by which he hopes to express Moore's real commitment, which White clearly realizes is just (G). But in that case, I don't understand what White has against Moore's account. Unless White considers (G) *itself* to be crazy, even where he's not at all confounding it with (H). Could that be it? I suppose it's largely because I really can't see why anyone would think that (G) itself is clearly false that I suspect that White rejects Moore's account at least in part because he is confounding (G) with (H).

There are places in the published literature where it at least seems that the fallacy of semantic descent is being committed against contextualism about knowledge.[24] But in conversation, in my philosophical travels, I believe that, even among philosophers of language and of mind, objections based on the fallacy of semantic descent form the single most common type of objection I've encountered. That is, the most common type of objection—at least as I'm inclined to divide types—is roughly of the form 'According to you …', or 'According to contextualism …', where the complaint is completed by filling in

---

[23] See, for instance, DeRose (1991: 582–4).

[24] Indeed, I'd count us as having witnessed an example in this chapter: I would count Yourgrau's use of his dialogue as an example of the fallacy. The contextualist is not committed to licensing B's crazy last line. That line is just an object-level analogue of some metalinguistic statements that, according to contextualism, B could appropriately and truthfully utter. Some other examples involve claims that contextualism implies statements of the types we've encountered in our Lewis passages 1–4.

some pretty implausible-sounding statement that I would in fact reject, that my views would lead me to reject, and which is such that the best explanation for why it's being attributed to me is that it's being confused with (or mistakenly taken to follow from) some higher-level statement which I really do accept.

Let me quickly give one quite common example that's a bit unlike the other examples of the fallacy that we've seen in this paper, in that it doesn't involve Lewis-like statements about knowledge 'vanishing' or 'going away', etc. I've frequently heard, occasionally from some *very* smart philosophers, roughly this: 'According to you, whether or not *Fred* knows depends on *our* context. But how can *Fred's* knowledge depend on *our* context?!' Now, I would never say,

(I) Whether or not Fred knows depends on our context.

In fact, my views would warn me against saying any such thing. What I would say is such things as,

(J) Whether we can truthfully assert various knowledge-attributing and knowledge-denying sentences about Fred depends (at least in part) on our context,

This is because the content of those assertions are context-dependent, according to me. And, using our semi-technical terminology of 'counts as knowing', I would say:

(K) Whether Fred counts as knowing in our context depends on what standards govern our context.

Now, maybe some of these objectors realize that I really reject (I) and are just using (I) as what they hope is an obvious expository shortcut to express my real view, (J) or (K) (where, for (K), 'counts as knowing' is used in our semi-technical sense). This is unlikely, though, because when I do tell such objectors that, like them, I reject (I), they're typically quite surprised by this. I have to suspect that such objectors are at least usually, and at least to some extent, committing the fallacy of semantic descent.

It is surprisingly and frustratingly hard to stamp out the negative influence of this type of objection to contextualism. This gives me cause to wonder whether I might be misunderstanding the objectors somehow. Maybe they have something more reasonable in mind which I've just been unable to discern? But I've tried very hard for a very long time to better understand, and haven't come up with anything. At this point, I think the best strategy is to make the charge that the fallacy is being committed, and see if anyone can then explain what, if anything, is going on instead. To do that, I'll focus on one

recent prominent example of what I take to be the fallacy being committed in the literature.

14. Dretske and the Fallacy of Semantic Descent?

So, here is the rightly renowned epistemologist Fred Dretske describing in print what he takes to be some of contextualism's 'bizarre' features:

In ordinary contexts in which knowledge claims are advanced, one knows (such things as) that there are cookies in the jar and (at the zoo) zebras in the pen. In such ordinary contexts one normally has all the evidence needed to know. One sees the cookies and the zebras and, given the absence of any special reason to doubt, this (for reasonably experienced adults) is good enough to know. So skepticism is false. Do these reasonably experienced adults also know, then, that they are not being deceived by some kind of fake? That the sensory evidence on which they base their judgments is not misleading? Well, yes (so closure is preserved), but—and here is the kicker—only as long as they don't seriously consider these heavyweight implications or say that they know them to be true. For once they think about them or say they know them to be true, the context in which knowledge claims are evaluated changes and knowledge evaporates. In this altered context (no longer an ordinary context) one doesn't know that one is not being deceived because new alternatives (that one *is* being deceived), possibilities one cannot evidentially eliminate, have been introduced. Therefore (closure forces one to say) one no longer knows that there are cookies in the jar and zebras in the pen. One gets to know about cookies and zebras only as long as one doesn't think about or claim to know what, according to closure, one has to know in order to know such mundane things. According to this way of dealing with skepticism, philosophers who spend time worrying about heavyweight implications (How do I know that I'm not dreaming? How do we know there is a material world?) are the most ignorant people in the world. Not only don't they know these heavyweight implications (maybe no one does), they don't (like everyone else) know the things (that there are cookies in the jar, zebras in the pen) that imply them. This, of course, includes almost everything they thought they knew.

This result, it seems to me, is pretty bizarre—*more* bizarre, in fact, that abandoning closure. It is a way of preserving closure for the heavyweight implications while abandoning its usefulness in acquiring knowledge of them. One knows (or is evidentially positioned to know) heavyweight implications so long as one doesn't think about them or say that one knows them to be true.   (Dretske 2005a: 18–19)[25]

---

[25] There might be *possible* views that Dretske's words here can accurately describe. That he intends to be describing contextualism as we're here construing it, and as it's held by Lewis, Cohen, and myself, seems to be indicated by a note he attaches nearby to the passage I quote (Dretske 2005a: 24 n. 3), and is also indicated by the facts that, just before the material I quote, Dretske introduces the view he is about to describe as a 'popular maneuver' (2005a: 18) among philosophers, and that when John Hawthorne

Note first that Dretske ascribes to contextualism just such commitments as I've been denying here (and have been denying for years). Dretske's contextualist holds, for instance, that:

(L) Our *knowledge evaporates* when we move into demanding contexts

(M) Believers know various ordinary propositions *only as long as they don't seriously consider* some of the 'heavyweight implications' of those beliefs such as that they're not being deceived in various ways.

(N) Philosophers who spend much time seriously considering such 'heavyweight implications' *are the most ignorant people in the world*!

Dretske's contextualist sounds a lot like Lewis. The first two of these three statements are very much of a piece with the Lewisisms, (1)–(4), that we have already discussed.

But note secondly that Dretske not only takes these commitments to be implausible in the extreme, but is pretty clearly presenting his claim that the contextualism has these 'bizarre' results as a reason to reject the view. Here he departs radically from Lewis, who knows better. On contextualism, no knowledge evaporates when we move into demanding contexts: one still has all the 'knowledge'-by-low-standards that one had before, and one never had the 'knowledge'-by-high-standards that one now lacks. And if it is the case that some philosophers spend more time than other people utilizing extremely high standards for knowledge that are almost never met, they are not thereby rendered more ignorant than other people any more than NBA scouts (and those who study skyscrapers and mountains) are rendered shorter (or are rendered more often short) than other people by the fact that they spend a lot of time using relatively hard to satisfy standards for 'tall' and relatively easy to satisfy standards for 'short'. One who tried to increase her knowledge by employing the ostrich-like strategy of scrupulously avoiding any conversation that employs high epistemic standards would utterly fail to protect anything that can rightly be called 'knowledge', for she would still lack all of the 'knowledge'-by-high-standards that she lacks in the case in which she's less careful about the conversational company she keeps.

What's going on here? Does Dretske fully realize that, for instance, contextualists don't really think that the philosophers in question are the most ignorant people in the world? Is he perhaps just 'bending the rules' a bit and

points out that Dretske has misconstrued the 'standard' version of contextualism (Hawthorne 2005: 39), Dretske's response to this complaint (Dretske 2005b: 44–5), to the extent that he does respond, contains no element of explaining that he never meant to be targeting such contextualist views.

using (N) as an 'expository shortcut' by which he hopes to express the *real* 'bizarre' implication of the contextualist's view? But what then is that? Perhaps something along the lines of:

> (O) Philosophers who spend much time seriously considering such 'heavy-weight implications' of their beliefs are the people in the world who most frequently use 'know(s)' in contexts in which it is governed by epistemic standards that they themselves do not meet!

But then I don't see why someone would take (O) to be 'bizarre'. (There are other candidates for the real feature of the contextualist view that Dretske might be using the likes of (N) to express, but none of the these candidates that are fair presentations of the contextualist's views seems much more bizarre than does (O).) So I suspect that, but am far from certain that, Dretske is committing the fallacy of semantic descent here.[26] Now, if what's really going on is that Dretske fully realizes that the contextualist doesn't actually accept anything like (L)–(N), and Dretske is only using the likes of (L)–(N) as 'expository shortcuts' by which he hopes to express the contextualist's real commitments, but Dretske does find those actual commitments—the likes of (O)—to themselves be 'bizarre', then no such fallacy is responsible for Dretske's own personal revulsion for contextualism. (In that case, at least by my lights, his 'bizarreness detector' is just wildly out of whack.[27]) But even so, he should stop using the likes of (L)–(N) to express the 'bizarre' features of the contextualist view to others, who might just think that contextualists really accept such things as (L)–(N).

In light of this sad situation, I think it unwise to engage in any rule-breaking expository shortcuts to express the contextualist's views and commitments.

---

[26] For those interested in digging deeper to figure out what Dretske's own thought is here, a potentially relevant passage to consider is Dretske (2005b: 45–6). For reasons similar to those I am pursuing here, Hawthorne makes the charge that in the passage of Dretske (2005a) that we've been considering here, Dretske betrays 'a flawed understanding of contextualism' (Hawthorne 2005: 39). In his response to Hawthorne (Dretske 2005b), Dretske mentions this charge of Hawthorne's at the top of (2005b: 45), and may be responding to it at (2005b: 45–6). I myself have great difficulties understanding Dretske's response, and, though he mentions Hawthorne's charge, I am not certain that he intends to there defend himself from the charge of misunderstanding contextualism. (He can be read as simply standing by his construal of contextualism, I think. He certainly does not admit that he got it wrong.) Some of what he writes may perhaps point to the interpretation that he finds the correct, higher-level statements of the contextualist view to themselves be bizarre. However, he continues to present the contextualist's views in what strike me to be badly inaccurate ways. So in the end, I'm very unsure about how to understand him here.

[27] Note that I wouldn't say this about someone who merely claimed to find the likes of (O) somewhat implausible. I would then just say that my initial intuitions about (O) differed from such a person. By my lights, such a person's 'implausibility detector' would seem to be misfiring here, but wouldn't seem to be 'wildly out of whack'.

If folks find the contextualist's actual commitments—like (J), above—to be wildly implausible, then, while I can't myself understand why they think this, I guess that's fair enough. But given what seems to be a general tendency to be misled by the fallacy of semantic descent, so far from actually using the rule-bending shortcuts, like (I), to express the contextualist's ideas, we should instead be explicitly warning folks to keep those misleading statements out of their minds as much as possible as they evaluate contextualism. If one does use the relevant rule-bending 'expository shortcuts' to express the contextualist's commitments, then, in light of all the (at least potential, and I suspect, often very real) confusion that surrounds these issues, one should explicitly state that one is doing just that, and clearly and forcefully warn readers not to hold it against contextualism that it sounds implausible when it's put in such terms.

7

Knowledge, Assertion, and Action

Contextualism vs. Subject-Sensitive Invariantism

In this chapter, we continue our critical comparison of contextualism with SSI. Sections 1–4 concern one of the main battlegrounds between SSI and contextualism: how these views handle some third-person attributions of knowledge. It's here that SSI faces what I consider its greatest problem. I will explain why these third-person attributions provide such a key test for deciding between the two views, and I will argue that, though it may appear that each view has third-person cases that it handles better than its rival, yielding something like a draw between the two views on this battleground, the reality is that only contextualism adequately handles all of the relevant cases, and that it thereby obtains a very serious advantage over SSI on this issue.

The rest of this chapter concerns the very important topic of the connections that knowledge might bear both to proper assertion and to various evaluations, predictions, and explanations of agents' actions. Accounting for such connections is perhaps the main driving force behind the recent fascination with SSI, and Hawthorne, in particular, has influentially argued against contextualism on the grounds that it problematically breaks the connections that seem to exist here. However, it turns out that contextualism does the better job here, or so I will argue. After taking a brief look at some of the ways that 'know(s)' is used in connection with evaluations, predictions, and explanations of actions in section 5, I then go on, in sections 6–10, to answer Hawthorne's objections to contextualism in this area. Then in sections 11–12, I go on the offensive and argue that there are important features of our use of 'knows' in connection with actions that only contextualism, and not SSI, can account for.

1. Contextualism, SSI, and First-Person Cases

Contextualists have long focused too much on first-person claims to knowledge (and first-person admissions that one does not know something), and too little

on third-person knowledge attributions, and I have been the chief offender here: First-person claims to know, and especially the Bank Cases with which that paper started (and which are also presented in section 1 of Chapter 1 here) were the stars of the show in (DeRose 1992). Though the contextualist thesis that I presented partly in response to those cases was phrased in terms of there being contextually variable truth-conditions for the general 'S knows that p', which is intended to have first-, second-, and third-person knowledge attributions as instances, and though I did later in (1992) discuss some third-person knowledge attributions, the Bank Cases were the cases in reference to which contextualism was introduced and most thoroughly discussed.

The reasons for this focus on first-person cases, at least in my own case, were, first, that some of the historical philosophers that were most influential on me had such a focus, and I was following their example.[1] Second, and most important, I was myself most interested in first-person uses of 'know' because I already had my eye on the knowledge account of assertion, and the support that it derives from clashes that crucially involve first-person denials of knowledge—clashes of the form 'P, but I don't know that p'. And, third, SSI wasn't on the scene yet, and so hadn't received the very able defenses that it has since enjoyed, and, though I had considered possible rivals to contextualism like SSI, I didn't find that possibility very threatening, since I thought it was easily refuted by third-person cases like those we're about to encounter.[2] Once views like SSI are seen as serious rivals to contextualism, third-person cases, which provide the crucial test between the two views, become very important.

Recalling the introduction of the views from Chapter 1 (especially section 11), there are important similarities between SSI's and contextualism's treatment of first-person cases like the Bank Cases. Most notably, both views will typically rule that in such cases, both the speaker/subject's positive claim to know in LOW and his admission that he does not know in HIGH are true, despite the fact that the speaker/subject is in an equally strong epistemic position in the two cases—he has the same grounds for his belief, is equally confident that his belief is true, etc. Here the two views join forces against classical invariantism, which insists that, because the speaker/subject (has a true belief and) is in an equally strong position in the two cases, one of those assertions—either his claim to know in LOW or his admission in HIGH that he doesn't know—must

---

[1] Descartes's *Meditations* are conducted primarily in the first person. In the twentieth century, Moore and Wittgenstein often discussed first-person claims to knowledge, leading Norman Malcolm to entitle his paper about their views 'Moore and Wittgenstein on the Sense of "I Know"' (Malcolm 1976).

[2] See n. 5, below.

be false. The classical invariantist, recall, holds that there is a single set of epistemic standards that constitute the truth-condition for all instances of 'S knows that p', including first-person claims of knowledge, regardless of the subject's 'practical' situation or the conversational context of the speaker. Since the speaker/subject is in an equally strong epistemic position in the two Bank Cases, he either meets those invariant standards, or he doesn't. If he meets them, then the admission that he doesn't know in HIGH is false. If he doesn't meet them, then his claim to know in LOW is false. There is no way that both assertions are true.³ As we have seen, the classical invariantist will typically try to account for the type of linguistic behavior exemplified by speakers in cases like the Bank Cases by claiming that varying epistemic standards constitute *warranted assertability conditions* for different uses of 'S knows that p', including the first-person claims of our two Bank Cases. Against this, SSI and contextualism both hold that the standard that the speaker/subject must meet in order to be able to *truthfully* claim to 'know' is different in our two Bank Cases. That's how, though he is in an equally strong position in the two cases, the speaker can meet the standard that governs LOW and fail to meet the standard in place in HIGH, and how his positive claim to know in LOW and his admission in HIGH that he doesn't know can both be true.

These similarities, however, should not tempt us to conclude that SSI and contextualism give basically the same analysis of first-person cases, and differ only when it comes to third-person attributions of knowledge. For despite the similarities noted above, SSI and contextualism actually provide very different analyses of first-person claims to know.

On SSI, the varying standards that comprise a truth-condition of 'I know that p' are sensitive to factors that attach to the speaker/subject *as the putative subject of knowledge*, rather than as the speaker of the knowledge attribution. That is, features of the speaker's/subject's practical situation determine a single epistemic standard that governs whether the subject himself *or any other speaker*, including those not talking with the subject, can truthfully say that the subject 'knows'. Truth-relevant features of the speaker's/subject's situation determine whether or not he meets those standards. If he does meet those standards, then, to employ a handy piece of semi-technical terminology, he *simply knows* that the bank is open on Saturdays: Any speaker, no matter their context, would be asserting a truth if they said he 'knows', and any speaker would be speaking falsely if they said he doesn't 'know'. On the other hand, if he doesn't meet

---

³ The classical invariantist may tolerate 'grey zones', where both 'I know' and 'I don't know' are neither determinately true nor determinately false. If so, she may hold that it can turn out that neither the speaker's claim to know in LOW nor his admission in HIGH that he doesn't know is false. Still, she will hold that it can't happen that both of those assertions are true.

the standards set by his own practical context, then the speaker/subject *simply doesn't know* that the bank is open on Saturdays, and anybody who says he does 'know' is asserting a falsehood, no matter their context.

The contextualist will likewise hold that the truth-relevant features of the subject/speaker's situation determine how strong an epistemic position he's in—and thereby determine which standards he does and which he does not meet. And the contextualist of course agrees that the subject/speaker's conversational context selects low standards as the relevant truth-condition in LOW and high standards in HIGH, yielding the result that the claim to 'knowledge' in LOW and the denial in HIGH are both true. But for the contextualist those are standards that attach to the subject/speaker *as the speaker* of the knowledge attribution. As such, they govern his own current talk about himself, but other speakers in different contexts can apply different standards to him, depending on their own conversational interests. In neither of the Bank Cases does our subject/speaker either *simply know* or *simply not know*, according to typical versions of contextualism.[4] In LOW, he meets the standards set by his own conversational context, so his claim to 'know' is true; however, different speakers in other contexts with different interests in him may apply higher standards to him and truthfully say that he doesn't 'know' that the bank is open. Similarly, in HIGH, our speaker/subject doesn't meet the higher standards set by his own context, and thus his admission that he doesn't 'know' is true; however, according to the contextualist, depending on their own conversational interests, other speakers in different contexts may apply lower standards to him, and truthfully say that he does 'know'.

This all suggests that to decide between these two accounts of our first-person cases, we should look to third-person cases. Since SSI holds that the subject's practical situation sets the standard for whether any speaker in any context can truthfully say that the subject 'knows', while contextualism says that the standard set by the subject's context is not necessarily binding on other speakers in other contexts, it's natural in adjudicating between the accounts to look to third-person attributions made in other contexts to see if the

---

[4] This is not to say that nobody ever simply knows or simply doesn't know anything, according to contextualism. Most obviously, subjects simply don't know propositions that aren't true or that they don't believe. And, if the 'floor' for allowable standards (see Chapter 1, section 6) is higher than that of mere true belief, subjects will simply not know even various true propositions that they believe if their epistemic position with respect to them is so weak that they don't meet even the lowest possible standards. In all these cases, subjects are such that no speaker in any context can truthfully say that they know. And if it could happen that a subject is in such a strong epistemic position with respect to a true proposition that she believes that she meets all allowable standards with respect to it, she would simply know that proposition. I imagine that it would be controversial among contextualists whether that's really possible, but contextualism in itself does not obviously rule out simple knowing.

speakers in those cases are constrained to use the standards set by the practical situation of the subject, as invariantism predicts, or are free to apply to those subjects different standards, depending on the speaker's own conversational interests, as contextualism would have it. So, in the next section, we look to third-person cases to adjudicate between SSI and contextualism. But we do so not because the two views analyze first-person cases in the same way. Rather, in addition to the importance in its own right of handling third-person cases correctly, investigating third-person attributions is also the natural way to determine which of the two views' very different handlings of first-person cases is correct.

2. Third-Person Cases that Vindicate Contextualism

When we turn to third-person cases, what we find is that contextualism is vindicated by the fact that speakers are not constrained to always apply to subjects the standards appropriate to the subjects' practical situation in the way that SSI predicts.[5]

This is shown by the third-person cases we have already considered. The third-person version of the Bank Cases that was presented in section 7 of Chapter 2 would also work here, but we will use the Thelma/Louise/Lena cases presented in section 2 of Chapter 1. The reader may wish to review those cases, though the following brief reminders may jog your memory. Recall that Thelma in LOW (at the tavern) is discussing whether Lena knows that John

---

[5] The argument against SSI that I'm about to give is one that, at least in a condensed form, I first published—as an argument against a wide group of possible views that would include SSI—at DeRose (1996b: 112), when SSI didn't yet have any actual proponents (at least that I knew or know of), and so just represented a possible view that someone might take. In (1996b), the argument was aimed at what I was there calling 'contextualism of the first type'—a type of view that would have current forms of SSI, as instances and which is not a form of 'contextualism' as we are here using the term. (As I was there construing it, 'contextualism of the first type' would also include views which, like SSI, had the subject's situation determine the epistemic standards, but which focused on the subject more as an enquirer rather than, as currently defended versions of SSI do, as a practical agent.) As I'm about to do here against SSI, I there wielded against that (somewhat broader) target a third-person attribution of knowledge in which speakers naturally and appropriately apply to a subject standards that are appropriate to the conversational interests the speakers are taking in that subject, rather than to the subject's own context. And, as I will do here a bit more expansively in section 4, I there more briefly explained that contextualism—or 'contextualism of the second type', as I was there calling it—can handle other third-person cases in which speakers apply to subjects epistemic standards that are appropriate to the subjects' situations: 'Of course, for certain purposes, and especially in cases where we're judging whether a subject's belief is justified, we may wish to evaluate her belief relative to standards set by features of her context. But there's nothing in the second version of contextualism to rule this out: among the many standards a speaker's context may select are those relevant to the subject's context' (1996b: 112). In (DeRose 2004a), I aimed this argument specifically at SSI.

was at work that day in connection with the matter of whether Lena will know that she has won the mere $2 bet she placed on the proposition that John was at work and thus will know to collect her winnings. Meanwhile, Louise in HIGH (talking with the police) is also discussing whether Lena knows that John was at work, but in Louise's case this is being discussed with an eye toward the issue of whether Lena can testify to that fact in a very high-stakes criminal matter. If SSI were true, then both Thelma in LOW (at the tavern) and Louise in HIGH (talking with the police) would in their use of 'know(s)' seek to apply to Lena the standards appropriate to Lena's own practical situation. However, as we observed, each of our two speakers, behaving just as normal speakers would, apply to Lena the standards that the speaker's own context calls for, Thelma applying to Lena the low standards appropriate to LOW's conversational purposes and saying that Lena 'knows', and Louise applying to Lena the high standards appropriate to HIGH, informing the police that Lena 'does not know' that John was at the office.

Here it might be suggested on SSI's behalf that whether Lena is aware of it or not, the fact that police had been considering her as a potential witness in a very important criminal investigation is part of her practical situation, and thus, at least with respect to the matter of whether John was at the office, very high standards in fact apply to her, and (to again employ our handy piece of terminology) she *simply doesn't know* that John was in. Louise is aware of the fact that the police have been considering Lena as a potential witness, and thus knows to apply to her the unusually high standards that Lena's practical situation calls for. That Thelma *mistakenly* applies low standards to Lena can be explained by the fact that Thelma is unaware of the highly charged practical situation Lena faces. I actually doubt that the proponents of SSI would accept this suggestion and the potential escape from their problem that it provides, since, based on how they tend to handle other cases, I wouldn't suppose that they would judge that the mere fact that the police, whom Lena will never actually encounter in relation to this matter, have been considering her as a potential witness would have the suggested drastic effect on the epistemic standards that apply to Lena. (I suspect that defenders of SSI would instead seek to escape their problem here in the way we'll look at in the next section—a way that involves denying the truth of Louise's denial that Lena knows.)

But, at any rate, this suggestion cannot withstand scrutiny, because so long as their conversational interests in the matter of whether Lena knows continue to diverge, Thelma and Louise will still (with apparent truth and propriety) describe Lena in these very different ways even if they both have the same information relevant to the matter of what Lena's practical situation is. Suppose

Louise calls Thelma on her cell phone, telling her of her encounter with the police, informing her that the police had been considering Lena as a potential witness, but also telling her that the police have now definitely decided not to even contact Lena (perhaps because they now know that Lena didn't actually see John). However, since they already have Louise there, they do want to interview Louise herself one more time before she leaves. It's as she's waiting for that final interview that Louise is taking the chance to call Thelma, to let her know that she might be getting back to her apartment a bit late and so might not be able to make the early morning jog they planned to take together the next day. After she hangs up the phone, Thelma rejoins the conversation at the tavern about the bet. When asked again by someone new if Lena knows that John was at work, and so will know to collect her winnings, Thelma will, of course, apply to Lena the low standards appropriate to LOW and answer affirmatively, as before. Meanwhile, in again talking to the police, Louise will again apply to Lena the high standards appropriate to HIGH, and will again quite naturally *deny* that Lena 'knows' that John was at the office. (And Louise will *definitely* not say anything like: 'Well, she didn't know that John was at work when you first asked me, because you were then considering using her as a witness, but now that you've ruled her out as a potential witness, she does know that', or 'Assuming you've definitely eliminated Lena as a potential witness, she knows that John was at work. But if and when you do consider her as a witness, she will no longer know that'!) Here, behaving just as normal speakers do and should, Thelma and Louise apply to Lena different standards, each speaker selecting the standards appropriate to the context of her own conversation, despite the fact that both are aware of the same facts that are relevant to the matter of what practical situation Lena is actually facing, even where that situation is construed broadly enough to include the intentions of the police toward her in this case in which Lena will never actually encounter the police.

As the above example shows, speakers apply to subjects, including far-away subjects that the speakers are not talking with, varying standards, depending on the speakers' own conversational interests. In fact, as we saw in section 7 of Chapter 2, where we converted the original, first-person Bank Cases into third-person cases, you can take almost any good pair of first-person cases, where SSI and contextualism will agree that the affirmation of knowledge in LOW and the denial of knowledge in HIGH are both true, and easily convert them to third-person cases where the contextualist's claims that both assertions are true are about as solid as the corresponding claims about the first-person cases—but with respect to which SSI must now deny that both assertions are true.

These third-person cases provide a powerful objection—to my thinking, a killer objection—to SSI, which predicts that our speakers in the high-stakes conversation will apply the lower standards to our subject. For SSI holds that the situation *of the subject* sets the standards that govern all attributions and denials of knowledge to that subject, and our speakers are aware of the facts that make the subject's context a low-standards one.[6] SSI does not necessarily predict that our speaker in HIGH will positively assert that the subject does 'know': If, for instance, a proponent of SSI also accepts the knowledge account of assertion, he will hold that since the speaker doesn't himself know that the bank will be open, the speaker also doesn't know, and therefore cannot assert, that the subject 'knows' this fact.[7] Presumably, then, we would expect the speaker to say something like that the subject *probably* 'knows'. But, in any case, on SSI, one thing we would definitely *not* expect of our speaker in HIGH is an outright assertion that the subject 'does not know': From the point of view of the speaker in HIGH, the subject is (at least probably) in a low-standards practical situation, and probably meets those relatively low standards that, according to SSI, govern any talk in any context about whether or not she 'knows'. But, as I've already noted, speakers in this HIGH situation actually do say precisely what SSI predicts they will not say. It is clear that speakers in situations like the one our speaker faces in the third-person Bank Case will often, and appropriately, apply their own, higher, standards to the subject: If the speaker is counting himself as not 'knowing', then, when he is speaking of the subject as a potential informant in the way we have been imagining, and this subject is no better positioned than is the speaker, the speaker will, with propriety and apparent truth, say that the subject, too, 'does not know'. Thus, against very strong appearances, and against the methodology of ordinary language that we have looked at, the defender of SSI will

---

[6] It's a bit artificial to suppose that our speaker has as much information as I've given him about the subject's conversational situation. I did so to make the argument simpler. In more realistic situations, where the speaker doesn't know what standards govern the subject's context, SSI makes somewhat more complicated, but equally false, predictions: It predicts that the speaker will be very agnostic about whether the absent subject 'knows', since he doesn't know what standards are appropriate to the subject's context. (In fact, since the standards the speaker applies are extraordinarily high, he should be highly dubious that these are the standards that govern whatever context the subject is in.) But that's not how speakers in fact behave in our speaker's situation. If the speaker knows that the subject doesn't meet the unusually high standards that govern the speaker's context, the speaker will say that the subject, like the speaker himself, 'does not know'.

[7] The proponent of SSI need not accept the knowledge account of assertion to make such a move. Similar moves could be made by SSI theorists who hold that there are varying standards for proper assertion, and that our speaker here is in a context governed by extraordinarily high standards for proper assertion, as well as for knowledge, even if they don't hold that the varying standards for these two are tied together in the very tight way specified by the knowledge account of assertion.

have to deny that the denial of knowledge in our high standards third-person case is true—or that the positive ascription of knowledge in the low-standards third-person case is true.[8] This will be an especially awkward stance for the proponent of SSI to take, since they themselves are committed to the truth of the key claims in both LOW and HIGH in the first-person cases, and it seems equally compelling to hold that the key claims in the third-person cases are both true.

3. The Projection Defense

Perhaps the best defense for SSI from these kinds of cases is to deny the truth of the denial of knowledge in HIGH by means of a projectivist strategy like the one that Hawthorne pursues (2004: 162–6). To explain this strategy, and then why it fails, consider again our high-standards third-person case involving Louise (who is talking with the police). According to the projectivist strategy, speakers like Louise *rightly* sense that they do not know (they do not meet the standards set by their own high-stakes context), and they then *mistakenly* project their own ignorance onto absent knowers like Lena (who does meet the standards set by her own, more ordinary, context, and thus, according to SSI, *simply knows* that John was at work). Hawthorne briefly appeals (2004: 164) to some work in psychology that may suggest some mechanisms that might explain our 'tendency to overproject our own lack of knowledge to others' (2004: 163), but does not want to be

---

[8] If the backer of SSI were to accept that the denial of knowledge in our high-standards third-person case is true, while still accepting the truth of the positive ascription of knowledge in the low-standards case, he will have to maintain that, though our speakers are far away from her, and their conversation will never affect her in any other way, the fact that these speakers are thinking of and speaking of the subject as a potential informant about a high-stakes matter in the high-standards case is part of that subject's 'practical context' and thus raises the standards for knowledge that govern whether that subject, or anybody else, whether or not they are talking to our speakers, can truthfully say that the subject 'knows' that the bank is open on Saturdays. That seems pretty implausible to me. However, in light of relatively recent developments in the philosophy of mind and language, we should be ready to be surprised by factors quite external to a speaker affecting the content of her thought and speech. Thus, we should seek further evidence against this possible move on behalf of SSI. That evidence is to be found in how the subject herself, or others talking about her, will describe her. Suppose the subject herself thinks it is quite likely that our far-away speakers are now applying very high standards to her in connection with some very serious matter. If that were enough to greatly inflate the standards that govern whether anybody could truthfully describe our subject as 'knowing', our speaker should then refuse to say that she 'knows'. After all, from her own point of view, such a claim would then be very likely false. But clearly there are many low-standards contexts where the subject's conversational purposes will call on her to simply apply to herself the low standards appropriate to her own very local context, ignoring what is or might be transpiring in far-away conversations.

committed here to any particular way of working out the details of such an account.⁹

But whatever particular psychological mechanism might be appealed to, it seems that a projectivist strategy like Hawthorne's, according to which speakers project their own ignorance onto other subjects, cannot provide the needed relief here for SSI. The problem with this projectivist strategy is that in the relevant cases speakers like Louise will deny that subjects like Lena 'know' *even when the speakers have no ignorance to project*. To see this, suppose that Louise *does* meet even the extraordinarily high standards that govern her context—she does clearly recall herself seeing, and even having a conversation with, John. So, in this revised case, Louise does know, and has claimed to know, that John was in the office. Still, if the police are wondering whether far-away Lena might *also* know (perhaps so that they can have two witnesses to that fact), Louise (who knows and takes herself to know, and thus has and realizes that she has no ignorance to project) will still deny that Lena knows if all Lena is going by is having heard the report and having seen John's hat.

To my above argument, Stanley responds that a form of projectivist strategy that's at least in the same ballpark as Hawthorne's could survive my objection by making the following modification.¹⁰ Instead of claiming that speakers in high-standards contexts like Louise mistakenly project their own *ignorance* onto subjects like Lena, the projectivist strategy could instead posit that speakers like Louise mistakenly project onto subjects like Lena the higher epistemic *standards* that are appropriate to the speaker's own practical situation. Mistakenly applying to Lena the standards that are appropriate to Louise's own highly charged situation, Louise will mistakenly judge that Lena does not know.

While the projectivist account Stanley suggests is *possible*, he provides, and I can see, little reason to actually accept it. Stanley does not explain how misleading tendencies of overprojection that we might display elsewhere can be applied to the case at hand to promote the thought that our tendency to deny that subjects 'know' when we are in such situations is a mistake, and it is extremely difficult, to say the least, to see how the cases from the psychological literature that Hawthorne appeals to could be convincingly used for Stanley's

---

⁹ As we will see in a longish passage from Hawthorne that I will soon quote, the mechanism Hawthorne has mostly in mind is the distorting influence of our use of an 'availability heuristic'. Nagel (n.d.) contains a good, critical discussion of the prospects for defending non-skeptical invariantism by appeal to such a psychological account.

¹⁰ Stanley (2005: 100–1). Stanley is responding to DeRose (2005b: 186–7), where I first raised the above objection to Hawthorne's strategy.

purpose here. Hawthorne appeals to the 'psychological literature on heuristics and biases', which he describes as follows:

Psychologists in that tradition emphasize the role played by the 'availability' heuristic as a distorting influence on our judgments of risk: in many cases, our estimation of the likelihood an event is affected by the ease with which we can recall or imagine it. So, for example, when a certain scenario is made vivid, the perceived risk of that scenario may rise dramatically. In this regard, it is a widely documented phenomenon that 'a recent disaster or a vivid film' may, as Slovic, Fischhoff, and Lichtenstein emphasize, 'seriously distort risk judgments' as, for example, when 'Recently experienced floods appear to set an upward bound to the size of loss with which managers believe they ought to be concerned.' They go on to note that a 'particularly important implication of the availability heuristic is that discussion of a low-probability hazard may increase its memorability and imaginability and hence its perceived riskiness.' [See also Johnson, Hershy, and Meszaros (2000), who found that people were willing to pay more for flight insurance that provided life insurance in case of terrorism than for flight insurance that provided that same amount of life insurance in case of death 'for any reason', remarking that 'events associated with "terrorism" … would be more vivid and available than events suggested by the inclusive phrase "any reason"' (p. 228).] Applied to the issue at hand, the availability heuristic may help to explain our tendency to skeptical overprojection. When certain non-knowledge-destroying counterpossibilities are made salient, we overestimate their real danger; as a result, we may find ourselves inclined to deny knowledge to others in cases where there is in fact no real danger of error.   (Hawthorne 2004: 164; the material in brackets above is Hawthorne's footnote)

While the connections between the very interesting behavior described above and the uses of 'know' in third-person cases like our original version of Louise's case (in which Louise has no more evidence than Lena has) are not very tight, and though Hawthorne did not offer enough help on how to apply these interesting psychological results to those uses of 'know' nor on why we should accept that application,[11] one can see how Louise's situation in the original version of our case might look similar enough to the situations of the subjects

---

[11] These problems become very pressing when we keep in mind that the defender of SSI needs to explain not only how we tend to make overly pessimistic judgements about the knowledge of others in the relevant cases (where speakers in a high-standards context negatively evaluate subjects in a low-stakes setting), but also, and simultaneously, why we are not in the same way inclined to overly pessimistic, *mistaken* misjudgements about our own knowledge, for in the relevant cases, the defender of SSI using a projectivist strategy as described above seeks to *endorse* the speaker's claim that she herself does not 'know', while rejecting as a mistake the speaker's claim that an absent third party similarly does not 'know'. It's this contrast that needs explanation. However, when Hawthorne gets to some of the details of the psychological results that he thinks make the strategy plausible, as he does in the longish passage quoted above, at least much of what we find would seem to impugn as overly pessimistic the relevant first-person judgements to the degree that they cast a shadow over the relevant judgements about other subjects.

described above to make it initially plausible to suppose Louise too may be engaged in somewhat similar 'skeptical overprojection': In that original case, it can seem plausible to suppose that Louise's highly charged situation might well be causing her to 'overestimate [the] real danger' that John might not have been at the office, which, it might be suggested with some plausibility, might lead her to draw a mistaken negative conclusion about whether Lena knows he was there. But the connection between the behavior described by Hawthorne above and the *revised* version of HIGH (in which Louise herself does meet even the extraordinarily high standards for knowledge) seems far too remote to offer Stanley support. In the revised case, Louise herself saw John at the office, and even had a conversation with him. From her point of view, there is no substantial risk (and, in all likelihood, as she'd be inclined to describe her situation, simply no risk at all) that John wasn't there. So she does not seem to be in any relevant way 'overestimating [the] real danger' that John wasn't there. And in fact she has claimed to know that he was there, despite the fact that she is in a highly charged, high-stakes context. So Louise's situation in our revised case seems very different indeed from those Hawthorne describes.

The crucial question about the revised case that currently faces us is: How will Louise describe Lena, who is in a low-stakes situation and who has only the limited evidence to go by? Absent any good reason to think otherwise, SSI predicts that Louise will apply to Lena the low standards relevant to Lena's situation, and say that Lena knows.[12] However, when we consider how 'know' is actually used, this prediction proves exactly wrong. Normal speakers in situations like Louise's will, appropriately and with apparent truth, apply to a subject like Lena the high standards appropriate to the speaker's own conversational purposes, and will say that the subject does not 'know', at least in cases like this, in which the distant subject is being considered as a potential informant to people in a high-stakes situation. It is a cost to SSI that it has to rule this very natural use to be a mistake. It does nothing to mitigate the cost involved simply to use the verb 'project', saying that speakers mistakenly *project* the standards appropriate to the speakers' own context onto absent subjects, rather than saying that speakers mistakenly *apply* the speakers' own standards to those subjects. Of course, appeals to the notion of projection could be helpful if they're accompanied by examples of clearly mistaken judgements people are led to make by 'projectivist' thinking which are similar enough to

---

[12] Because Louise herself does know, even by the extraordinarily high standards, that John was present in this revised version of HIGH, we do not face the complication we discussed that arises in the original version of HIGH: that because Louise doesn't herself know that John was in, she is not in a position to assert that anyone else knows this either.

relevant uses of 'know' to substantially support the thought that those uses are mistaken as well. And it is conceivable that certain psychological results reveal, or will reveal, certain patterns of mistaken judgements that people tend to make which can provide SSI with the relief it needs here, because the judgements speakers issue in the relevant third-person cases fall into patterns of what are clearly mistakes that we tend to make to a degree sufficient to render it plausible that those third-person judgements about what is and is not known are also mistakes. But it's also possible, for all that's been shown, that psychological results reveal or will reveal that SSI is in even deeper trouble than we've so far discerned.[13] But at this point, it seems that SSI simply issues a mistaken prediction about the very cases that are crucial to deciding between it and contextualism, and there's nothing we've seen in the projection defenses we've looked at to trouble that verdict.

4. Other Third-Person Cases: A Big, Ugly Tie?

For reasons that have become apparent, I think that SSI's fumbling of the third-person uses of 'knows' that we've looked at provides a quite decisive objection against the view. It is decisive because, first, given the nature of the two views we are considering, third-person cases like the ones we have just considered naturally provide key tests to discern which view is right. Thus, SSI's failure here is lethal to its hopes of matching contextualism's performance on this crucial test. And, second, the premises on which this objection to SSI is based (concerning the truth of the third-person claims in both LOW and HIGH) are about as solid as are the similar judgements concerning first-person cases, where SSI agrees with contextualism that both assertions are true.

Defenders of SSI of course see things differently. This is no doubt partly because they don't think the problem is as damning as I think it is, but also partly because they think SSI's rivals, including contextualism, face equally daunting challenges. It's in this spirit that Hawthorne writes: 'This is not to say, of course, that there are no counterintuitive consequences to this version of sensitive invariantism. As far as I can see, every candidate story about our puzzle has counterintuitive results. This is no exception.'[14]

---

[13] Nagel (2008, n.d.) contain good discussions of the potential bearing of recent psychological results on the contest between contextualist and various invariantist accounts of knowledge attributions, and can also serve as good guides to philosophers to what are some of the relevant psychological results.

[14] Hawthorne (2004: 162). I'm counting Hawthorne as a defender of SSI, and I think he can be accurately counted as such, since he does lean in that direction. But the reader should be aware that

I believe that none of the recent problems alleged by Hawthorne and by Stanley to afflict contextualism are even close to being as powerful as the objection SSI faces in virtue of its mistreatment of third-person attributions and denials of knowledge, though this is based in several cases on answers to those objections I believe are available—many of which we've already seen, especially in Chapter 5. But there is a potential objection to contextualism that can appear to be extremely strong, that is directly related to the objection to SSI that we've just encountered, and that can seem to balance off the relative disadvantage that SSI faces in the cases we've just looked at, but that can be answered by an important clarification of contextualism. So we will discuss that objection here.

In the high-standards third-person cases I have discussed here in my attack against SSI, the speaker applies extraordinarily high standards, appropriate to the speaker's own high-stakes context, to an absent subject who is not herself in such a high-stakes context (nor does the speaker take the subject to be in such a context). And we speakers often do apply standards appropriate to our own contexts to far-away subjects in quite different contexts. Among the situations in which we will naturally do this are contexts in which we are discussing these absent subjects as potential informants to us—as our speakers are doing in the third-person cases I have used.

But there are other conversational situations—situations that advocates of SSI rightly pay a lot of attention to, and that we will pay a lot of attention to in the remainder of this chapter—where different conversational purposes are in play, and in which speakers will naturally apply to absent subjects standards that are appropriate to the practical contexts *of those subjects*. This often happens when, for example, the speakers are discussing practical decisions those distant subjects face: e.g. 'She should do that only if she knows that...' SSI handles such cases nicely. Indeed, these cases may appear to strongly favor SSI over contextualism, since, according to SSI, the subject's context sets the standards, while, as Hawthorne writes, according to contextualism, 'It is always the ascriber's standards that call the semantic shots, so to speak' (2004: 59). In fact, this can seem to be an advantage for SSI that balances the advantage I've claimed for contextualism: Sometimes it seems that the speaker's context sets the standards for third-person uses of 'knows'; sometimes it seems the subject's context sets the standards. SSI doesn't work well on the first kind of case; contextualism on the second. It's a wash.

Hawthorne does not have a strongly settled view on what is the right theory here; on this, see especially (2004: 187–8).

If that *were* the situation we faced, I would consider that to be a tie that renders both sides losers, since I take failures of this magnitude very seriously.

Fortunately, that is not the situation we face, for contextualism can in fact easily handle the second type of case. Two points are vital here. First, there is nothing in contextualism to prevent a speaker's context from selecting epistemic standards appropriate to the practical situation of the subject being talked about, even when the subject being discussed is no party to the speaker's conversation—which is good, because speakers often do select such standards when their conversational purposes call for it! And second, not only does contextualism *allow* for the *possibility* that the speakers' context will select standards appropriate to the practical situation of the subject, but, in the relevant cases, it would actually lead us to *expect* that the speakers' context will select such standards, because the *speakers' own conversational purposes* call for such subject-appropriate standards in the cases in question. On contextualism, the speaker's context does call the shots. Hawthorne is right about that. But sometimes speakers' own conversational purposes call for employing standards that are appropriate to the practical situation of the far-away subjects they are discussing, and so the shot that the speakers' context calls can be, and often quite naturally will be, to invoke the standards appropriate to the practical situation faced by the subject being discussed. And one type of conversational situation in which it will be quite natural for speakers to (and perhaps even strange for them not to) employ standards appropriate to a subject's situation is contexts in which the speakers are discussing whether or not the subject 'knows' in connection with evaluating a practical decision or action (including, it's worth adding, the action of asserting something) that the subject faces (or faced, will face, could have faced, etc.): 'She should do A only if she knows that p', 'Well, if he knew that p, he could have responsibly done A', 'She should not have said that p, because she didn't know that it was true'. *In such a context*, where whether the subject 'knows' is being discussed for such a purpose, it would often seem strange to employ epistemic standards not appropriate to the practical situation faced by the subject. It's an all-important feature of the speaker's own context that makes it natural for standards appropriate to the subject's practical situation to be invoked—namely, that the subject is being discussed precisely in connection with such an evaluation, prediction, or explanation of a practical decision of hers. So, supposing contextualism, it's no surprise at all that when such interests are in play, the speaker's own context will select standards appropriate to the practical situation faced by the subject.

But, as is equally evident, and as we have already seen, there are other, very different contexts in which we have quite different reasons for discussing

whether a subject 'knows', and in which it is perfectly appropriate to, and in which speakers actually do, apply to subjects standards quite different from those that the subject's own practical situation calls for (and not because the speakers are mistaken about the nature of the subject's context). And, disastrously, SSI cannot handle this equally evident fact, since SSI has the subject's situation set the standards that govern any speaker's description of that subject, whatever the speaker's context. The 'big ugly tie' is therefore avoided: One of the two views under consideration can handle both types of third-person cases.

5. Some Uses of 'Know(s)' in Evaluating, Explaining, and Predicting Actions and Assertions

The realization that a speaker's context may 'call the shots' by selecting standards appropriate to the practical situation of the subject being talked about, and that often the speaker's context actually will call such a subject-appropriate shot where what's being discussed is how the subject should act (or should have acted), is also key to answering one of the main criticisms Hawthorne raises against contextualism: the attack he develops at (2004: 85–91) that contextualism implausibly breaks some important ties that knowledge bears to assertion and to action. In fact, I think it can be shown that contextualism can accommodate as tight a tie here as the data support, and in the final analysis handles the ties that exist between knowledge and action better than does SSI. But we will have to work up to that point. We will start in this section by taking a quick and rather breezy look at some of the linguistic data to be explained concerning the connections that knowledge bears to proper assertability and to action.

It certainly is true that our evaluations of whether or not subjects (ourselves or others) 'know' propositions to be the case are often made in connection with, and often in support of, evaluations we make about those agents' actions (and potential actions)—including their verbal actions of flat-out asserting things. The judgements being supported can be evaluations of the morality or rationality of the actions involved, or can be any of a number of other types of evaluations. As noted in the above section, it is very common for us to say such things as 'She should do A only if she knows that p', 'It was irresponsible of him to A, because he didn't know that p', 'She didn't act foolishly in doing A, because she knew that p', etc. Also, whether a subject does or does not know a particular proposition is often cited in support of a judgement,

prediction, or explanation of what the subject will in fact do: Given the right circumstances, it is utterly natural for a speaker to say something like this: 'No, I don't think Bill will be late. He probably won't wait in the long lines at the bank, because he knows that the bank is open on Saturdays.' Or: 'Bill probably will be late. He's likely to wait in the long line at the bank because he doesn't know that it's open on Saturdays.' As such uses certainly seem to show, many of our evaluations, explanations, and predictions of agents' actions depend on whether they 'know' certain propositions to be true. (It also seems that such uses of 'know(s)' are extremely important, and that epistemologists should strive to make sense of the connections here.)

What's more, to cite a fact that proponents of SSI will certainly accept, and so does not require defense in this context (though it's pretty clearly true in any case), making sense of the range of the natural uses in question requires some appeal to varying epistemic standards. For, when describing an agent who faces a low-stakes situation, for instance, we might defend her behavior as in some way all right because she 'knew' that such-and-such was the case, where that agent was in only a moderately good epistemic position with respect to the true proposition in question, while we might condemn (as irresponsible, for example) the action of an agent in a highly charged situation because he *didn't* 'know' a certain proposition to be the case (even if that proposition is in fact true), where this latter agent is in the same moderately strong epistemic position as is the low-stakes agent whose behavior we praise. How strong an epistemic position an agent must be in to merit certain evaluations of their behavior (or potential behavior) varies with their circumstances, and we often accommodate that by saying that they do or do not 'know' things to be the case, where our use of 'know' is governed by different standards on different occasions.

Contextualism, of course, allows for the possibility of such uses of 'know(s)', with varying standards. Where a speaker's context selects epistemic standards such that a subject's 'knowing' or not 'knowing' some proposition according to those standards supports a certain evaluation of a subject's behavior, then speakers can profitably cite that the agent does or does not 'know' that the relevant proposition is true as a good reason for that evaluation of the agent's action or assertion. What's more, contextualism should not lead us to expect that it will only be in rare cases of lucky convergence that speakers' contexts select epistemic standards that are in that way appropriate to the practical situations faced by the subjects they are talking about. As I noted in the previous section, in contexts in which speakers are discussing whether or not a subject 'knows' something in connection with evaluating a practical decision that the subject faces (or faced, will face, could have faced), the speaker's own conversational purposes often make it natural for them to (and perhaps even

strange or wrong for them not to) employ epistemic standards appropriate to the subject's practical situation, and thus, supposing contextualism, we should expect such subject-appropriate standards to be selected.

6. Hawthorne's Charges that Contextualism Breaks the Connections that Knowledge Bears to Assertion and to Practical Reasoning

So, why might it be thought that contextualism problematically breaks the ties that knowledge seems to bear to assertion and to practical reasoning? The conclusion most obviously supported by the data noted in the above section is that assertions about whether agents do or don't 'know' various propositions often have a content that makes them suitable for supporting various evaluations and predictions of their behavior. But Hawthorne apparently thinks that, going beyond such fairly loose connections, the data support certain principles connecting knowledge with assertion and with rational action. In the case of the knowledge–assertion connection, Hawthorne thinks there are good grounds to accept what we have been calling the knowledge account of assertion: that a speaker is in a position to assert some proposition iff she knows it to be the case. And with regard to the connection between knowledge and action, Hawthorne and Stanley think the data support the following principle:

> *The Reason-Knowledge Principle (RKP):* Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting iff you know that p.[15]

As we saw in Chapter 3, I certainly agree with Hawthorne that there are good grounds for accepting the knowledge account of assertion, and I myself accept it—at least in the relativized form we've labeled KAA-R. And while, for reasons I'll give in section 8, I'm not similarly sold on anything like the Reason-Knowledge Principle, the contextualist who is more impressed than I am by the support for such a principle could accept a similarly relativized

---

[15] This name for the principle and this exact formulation of it come from Hawthorne and Stanley (forthcoming). A 'first pass' at trying to articulate such a principle occurs in Hawthorne (2004: 30). The exact account of what it is for the choice of an agent to be 'p-dependent' will not be important to our discussion. The purpose of the clause is to rule out cases where whether p is true is irrelevant to the choice in question. In section 10, below, I will suggest what seems to me a better way of formulating such principles. My suggestion will be made with respect to a 'one-way' version of such a principle, but can be applied just as well to 'two-way,' or biconditional, principles like RKP.

version of that principle as well, as we will see a bit later. So what's the problem? Hawthorne's main complaint seems to be that the contextualist is problematically committed to endorsing certain conjunctions that 'seem strange' and 'sound very odd' (2004: 87)—where these conjunctions violate the simple versions of the principles in question, but not the relativizations that the contextualist must accept if he is to back some version of the principles. A secondary worry Hawthorne seems to have is that accepting the contextualist's relativized versions of these principles doesn't amount to accepting the principles in question in a full-blooded enough way. We'll start with his main complaint.

7. Contextualism and Some Strange Sentences Concerning Knowledge and Assertability

Having just cited some examples in which we naturally cite whether agents do or don't know things in support of evaluations of their actions or assertions (2004: 85), Hawthorne begins his attack by identifying what he takes to be the root cause of this trouble for contextualism:

It is natural, then, to think that there is some deep association between facts about knowledge and facts concerning the propriety of assertion and practical reasoning. What I want to draw attention to here is that, owing to the purported ascriber-dependence of 'know', contextualism seems to disconnect facts about knowledge from these normative facts. And this is because the relevant normative facts do *not* seem to be ascriber-dependent.   (2004: 86)

Whether it is proper for subjects to assert propositions or to use them in various ways in practical reasoning depends only upon those subjects' own circumstances. This means that the only forms of KAA and of RKP that the contextualist can accept will be ones that are relativized to make appropriate assertion or use of reasons depend on whether the subject knows propositions *according to the standards appropriate to the subject's own practical situation*. Thus, as Hawthorne correctly notes in the last sentence of the above quotation, sentences evaluating the propriety of this behavior don't display the kind of *ascriber*-dependence that, according to contextualism, knowledge-ascribing sentences display.[16] It's this discrepancy, according to Hawthorne, that commits

---

[16] Such sentences may well display *some* forms of context-sensitivity. Hawthorne need not deny that. All he needs here is that these evaluations aren't sensitive to varying epistemic standards that are sensitive to the ascriber's context in the way that contextualism says that knowledge attributions are.

the contextualist to endorsing the assertions of certain sentences that can sound very odd. Hawthorne produces strange sentences involving both assertion and practical reasoning, but since he works out his attack most thoroughly and convincingly in the case of the knowledge–assertion connection, we will start by considering strange sentences involving proper assertion. The defense I give against Hawthorne here will carry over to his charge that contextualism is committed to some strange sentences involving knowledge and practical reasoning.

We will begin by looking at some other odd sentences concerning knowledge and assertion that are based in key ways on Hawthorne's own, but that are modified to be about particular assertions of particular agents. Addressing these modified sentences will nicely set up our discussion of the more general sentences on which Hawthorne's attack is actually based.

So, it can seem that even the contextualist who accepts only the relativized knowledge account of assertion (as I do) is committed to endorsing assertions of sentences like the following, in certain circumstances:

(P5) Louise knows that John was at work today, but she ought not to assert that he was because her epistemic position is not strong enough with respect to that fact.

(P6) Thelma flat-out asserted that John was at work today though she did not know that to be the case, but she is not subject to criticism for doing so.

What agents like Louise and Thelma are in a position to properly assert depends only on their own circumstances. Thus, as Hawthorne would correctly claim, if Thelma meets the standards for knowledge and proper assertion set by her context, but Louise does not meet the standards determined by her context, then the second half of each of (P5) and (P6)—by which I mean the portion that, in each case, follows the word 'but'—is *simply true*: it is true to assert, no matter the speaker's context. But given our claims about how knowledge attributions are ascriber-dependent, we contextualists seem committed to endorsing the first halves of (P5) and (P6) (the part of each that precedes the word 'but') in certain contexts, as well. In the case of (P5), suppose that, though Louise does not meet the standards for knowledge and for proper assertion in place in her own context, she does meet the lower standards for knowledge that govern some other speaker's context. Then, it seems, the contextualist would counsel that other speaker to assert the first half of (P5). And since the second half of (P5) is *simply true*, it can seem that the contextualist would counsel such a speaker to assert *all* of (P5). And an analogous case can be made for a contextualist endorsement of the assertion of

(P6) in a case where a speaker is in a context governed by high standards for knowledge that Thelma does not meet, even though Thelma does meet the lower standards for knowledge and for proper assertion set by Thelma's own context.

While (P5) and (P6) can sound odd or strange, I should register that, at least to me, they don't sound awful. And they certainly don't sound as bad as, say, some of the sentences we discussed in Chapter 3, such as Moore's example: 'Dogs bark, but I don't know that they do' (1962: 277). If it turned out that contextualism had to rule that (P5) and (P6) can express truths in some contexts, that *in itself* wouldn't be such a worry: After all, as we all know, even Moore's above truly awful-sounding sentence can be truthfully (if inappropriately) asserted—that's part of what's puzzling about Moore's example. Still, the oddness of (P5) and (P6) should be accounted for, especially in situations, if there are any, in which they are held to express truths. Thus, that contextualism seems to be committed to there being contexts in which these sentences express truths places a special burden on the contextualist to account for why they can seem so implausible, though people may differ over just how bad these sentences do sound.

We have already seen the essence of how the contextualist can account for the wrongness of assertions of the likes of (P5) and (P6). As we have noted, where a speaker is evaluating whether a subject 'knows' something to be the case in connection with such an evaluation of the propriety of that subject's own assertion, the speaker's own conversational purposes will call for her to employ standards appropriate to those evaluations of the subject's behavior. On this account, the contextualist agrees that (P5) and (P6) can be true relative to some epistemic standards, but can still account for why *asserting* these sentences will be problematic.

This account can take several different forms. Which form to adopt will depend in part on just how problematic (and how invariably problematic) one finds the likes of (P5) and (P6). First, those contextualists who find that these sentences grate on their ears very badly in just about any imaginable situation can maintain that, because they involve citing facts about whether subjects 'know' things in support of such evaluations of those subjects' assertions, an assertion of (P5) or of (P6) will *invariably* cause the epistemic standards relevant to the evaluation of the subject's assertion that each involves to govern the speaker's own assertion. To work this out in the case of (P5), suppose that the standards for warranted assertion and for knowledge that govern a certain speaker's context are currently low enough that Louise's belief that John was at work meets those standards, even though Louise's belief doesn't meet the higher standards of Louise's own context. This speaker, then,

initially finds herself in the very type of situation in which, according to the contextualist, (P5) is true relative to the standards that govern her context. Still, on this account, it will be wrong for her to assert (P5), for the very act of asserting (P5) will invoke the higher standards relative to which the beginning of (P5) is false. For analogous reasons, if a speaker starts in a context relative to the standards of which (P6) is true, she will be unable to truthfully assert (P6), because such an assertion will bring into play lower standards relative to which Thelma does count as knowing and the beginning of (P6) is false.

But the contextualist doesn't have to insist that assertions of the likes of (P5) and (P6) invariably invoke the standards relative to which their first halves are *false*. Those contextualists who find these sentences to be rather odd, though perhaps not gratingly awful, might prefer a second version of our account, according to which an assertion of such a sentence creates a pressure on the standards of a speaker's context strong enough to guarantee that the first half of the assertion *will not be true*. Again working this out in the case of (P5) (and this time leaving the application to (P6) completely to the reader), suppose again that a speaker starts in a relatively low-standards context relative to which (P5) is true. But now suppose that there is some pressure to keep the standards of the context down to a level relative to which (P5) is true. We will remain open as to the source of this pressure, but presumably it will involve conversational purposes other than the evaluation of the propriety of Louise's assertion that remain very pronounced even after our speaker asserts (P5). Perhaps all the participants in the conversation, including the one who asserts (P5), continue to talk in ways that indicate that the standards remain at the relatively low levels (at which (P5) is true), with nothing, other than the assertion of (P5), to count in favor of the standards rising to the level relevant to evaluating Thelma's assertion. The contextualist who utilizes the second version of the account does not claim that the assertion of (P5), in such a context, necessarily results in the higher standards relative to which the first half of (P5) is false being successfully installed (however briefly). Instead, she claims that in such cases, (P5) may simply go truth-value-less, because it is indeterminate what standards govern it. The assertion of (P5) very strongly calls for the invocation of the high standards relevant to its evaluation of Thelma's assertion, and this call is forceful enough to guarantee that (P5) will not determinately be governed by other standards. But perhaps in some rather defective contexts, the assertion of (P5), and the push that assertion provides toward the high standards relevant to the evaluation it contains in its second half, will coexist with a continuing strong conversational pressure toward lower standards. In such strange examples, it may be simply indeterminate which standards govern

the assertion of (P5). By contrast with our first version of the account, on this second version, the assertion of (P5) need not be false. Still, the second version of the account explains the wrongness of assertions of (P5) (and, analogously, of (P6)), because, on it, as on the first version, an assertion of (P5) cannot be true.

But the contextualist doesn't even have to go that far to explain the oddity of the likes of (P5) and (P6). On the third, and probably most plausible, version of our account, the contextualist can allow for the possibility that assertions of the likes of (P5) and (P6) can in some cases be true, but can still account for the oddity of them, because even such true assertions will at least often be problematic, badly misleading, and will constitute cases of speakers choosing to say the wrong thing. On this version of the account, the force that assertions of the likes of (P5) and (P6) exert toward the standards relevant to the evaluations contained in their second halves does not secure the wrongness of the assertions in virtue of that force's relative power to overcome, or at least neutralize, any opposing conversational force it might possibly run up against. Rather, it is claimed that even if there might be settings in which assertions of (P5) and (P6) are determinately governed by standards other than those relevant to the evaluations contained in their second halves, those evaluations demand the use of the standards that are relevant to those evaluations strongly enough that it is at least often wrong to assert the likes of (P5) and (P6) in such a setting. It seems more than reasonable to suppose that you should not cite facts concerning whether a subject 'knows' something in such close connection with an evaluation of the propriety of an assertion of hers when the fact you would then cite does not support the evaluation you are making and can be easily confused with a proposition that would support the evaluation. Some who adopt this third account may even hold that in some conversational situations, it may remain clear enough through an assertion of the likes of (P5) and (P6) that standards different from those of the subjects being discussed govern the speakers' context that an assertion of the likes of (P5) or (P6) may be appropriate as well as true. At any rate, in the relevant situations in which an assertion of the likes of (P5) or (P6) is to some degree wrong, despite perhaps being true, it can be held that speakers who might find themselves in such settings and who wish to *correctly* express the truths that can only improperly be expressed by assertions of (P5) and (P6) will have to express those truths with other, more complex sentences—perhaps something along the lines of the following, which, though complicated, closely resemble the types of sensible clarifications speakers often have to resort to when resolving with one another various claims that have been made using other context-sensitive terms (as, for example, when they are questioned about the propriety of claims they have

made in various contexts about whether various people are 'tall', where those claims are in surface tension with one another):

(P5r) Louise counts as knowing that John was at work today by the standards we are using in our current conversation, but she ought not to assert that he was at work because her epistemic position with respect to that fact is not strong enough to meet the higher standards for knowledge in play in her own conversation.

(P6r) Thelma flat-out asserted that John was at work today though she does not count as having known that to be the case by the high standards for knowledge that we are currently using, but she is not thereby subject to criticism, because she did meet the lower standards for knowledge that governed the more casual conversation she was engaged in at the time.

It is perhaps in part because (P5) and (P6) don't strike me as being all that problematic that I find the third of the above accounts to be the most plausible, and, indeed, find that third account to be quite adequate for explaining the degree and nature of the wrongness that I sense with those sentences. Those who find the sentence more deeply troubling may prefer one of the other accounts.

Our above contextualist account of the wrongness of (P5) and (P6) can be fairly straightforwardly applied to the conjunctions Hawthorne actually bases his attack on. Hawthorne utilizes the following sentences, which are not about particular agents, but are rather general statements:

(H5) There are things people know but ought not to assert because their epistemic position is not strong enough with respect to those things.

(H6) People often flat-out assert things that they do not know to be true but are not thereby subject to criticism.[17]

Hawthorne notes that these 'seem strange' and 'sound very odd' (2004: 87). Yet, it can seem that the contextualist is committed to endorsing the assertion of each of these in certain circumstances, for reasons similar to those that account for contextualism's apparent endorsement of (P5) and (P6): A speaker in a relatively low-standards context would seem, on contextualism, to be able to happily assert (H5), since (H5) applies, among others, to people who are in high-standards contexts, and contextualism would seem to counsel speakers in relatively high-standards contexts to assert (H6), because the people they would

---

[17] (Hawthorne 2004: 87). Hawthorne simply labels these (5) and (6).

thereby be describing would include some who are in low-standards contexts in which they're speaking in an unobjectionable manner in asserting things that they don't count as knowing by the higher standards that we are supposing govern the speaker's assertion of (H6). In reference to these sentences (among some others), Hawthorne concludes: '[T]here are a variety of sentences that seem intuitively incorrect … but that, by the lights of contextualist semantics, come out true. That is surely a cost of the view' (2004: 89). As with (P5) and (P6), while I agree that these sentences can sound somewhat odd, they don't sound to me all that terrible (and Hawthorne's own sentences sound a bit less problematic than do (P5) and (P6)), and I wouldn't take it to be a great worry in itself for contextualism if it had to rule that there are contexts in which these sentences express truths, though I do think their oddness has to be explained.

But we are now prepared to fairly readily see why, on contextualism, it can be somehow wrong to assert the likes of (H5) and (H6). Again, the key insight is that it is to be expected, given contextualism, that uses of 'know' that are made in connection with such evaluations of the propriety of subjects' assertions will be governed by the standards that, given the situations of the subjects being described, are appropriate to those evaluations. Applied to general statements like the ones currently under consideration, this insight should lead us to expect that the likes of (H5) and (H6) will be somehow wrong to assert if they are governed by particular standards that diverge from those appropriate to judging the assertions of some of the people to which these general statements apply.

Again, the exact nature of the wrongness involved is open to negotiation. A close analogue of the first version of our contextualist account of the wrongness of (P5) and (P6) seems not to be available to account for the wrongness of the more general (H5) and (H6). The contextualist cannot suppose that *the* set of standards relevant to the evaluations of the assertions of the subjects involved will be invoked by an assertion of (H5) and (H6), because (H5) and (H6) apply to many subjects, relative to the evaluation of which many different standards are relevant. But fairly straightforward adaptations of the second and third versions of our contextualist accounts of the wrongness of (P5) and (P6) seem applicable to (H5) and (H6). Spelling this adaptation out in the case of the third and most plausible version of the account yields the following. Even if there are situations in which a speaker could assert the general (H5) or (H6) and his assertion would be governed by particular epistemic standards that are inappropriate for the evaluation of some of the agents' assertions to which the general statement applies and relative to which (H5) or (H6) are true, it is nonetheless at least often wrong for a speaker to assert (H5) or (H6) in such

a situation, precisely because such an assertion would involve citing whether agents 'know' things in support of an evaluation of the propriety of their assertions, and so should not be governed by epistemic standards inappropriate to evaluation of the assertions of (some of) those agents. Speakers who might find themselves in such settings and who wish to *correctly* express the truths that can only improperly be expressed by assertions of (H5) and (H6) will have to express those truths with other, more complicated sentences—perhaps something along the lines of the following:

(H5r) It often happens that people know things according to the standards we are using in our current conversation but that they ought not to assert those things because their epistemic position is not strong enough with respect to those things to meet the standards for knowledge that are in play in their own conversations.

(H6r) People often flat-out assert things that they do not count as knowing by the standards we are currently using, but, because they do meet the standards for knowledge governing their own conversations, they are not thereby subject to criticism.

To my thinking, this third explanation more than adequately accounts for the problematic nature of (H5) and (H6)—though, again, that may in part be due to the fact that those sentences never sounded all that terrible to me to begin with. Those who find the sentences more damningly grating can instead opt for the second account.

8. Contextualism and Hawthorne's Strange Sentence Concerning Knowledge and Practical Reasoning

For reasons analogous to those we saw toward the beginning of the previous section, it can appear that contextualists are committed to endorsing the use in some contexts of the likes of these analogues of (P5) and (P6), adjusted to concern practical reasoning rather than proper assertion:

(RP5) Louise knows that John was at work today, but she ought not to use that fact in her practical reasoning, because her epistemic position is not strong enough with respect to that fact.

(RP6) Thelma used the fact that John was at work today in her practical reasoning, though she did not know that to be the case, but she is not subject to criticism for doing so.

The reader can also construct analogues of the more general statements (H5) and (H6), to which the contextualist will appear to be committed. These new sentences concerning practical reasoning intuitively strike me as being even less problematic than their analogues concerning proper assertion. That may be in part due to the fact that these new sentences concern the fairly theoretical notion of what propositions can be properly used in one's practical reasoning. At any rate, the accounts given in the last section for why the sentences involving proper assertion can seem wrong can be easily adjusted to apply to these new sentences, and the third account would seem to me to be more than adequate for accounting for the problematic nature of these sentences concerning knowledge and practical reasoning.[18]

However, in his attack on contextualism here, Hawthorne produces a sentence concerning knowledge and practical reasoning that sounds considerably

---

[18] At their (2007: 576), Fantl and McGrath report that in a manuscript they have seen, Stewart Cohen accounts for the wrongness of conjunctions of the same basic type as we are now dealing with in what looks (from Fantl and McGrath's description) to be something like this first of the three ways I have put forward here. Fantl and McGrath then go on to critique the maneuver (2007: 576–81). Though they are specifically aiming their criticisms at something like the first of the three accounts I'm considering, and that first account is not the one I think is the most plausible, and though some of their criticisms seem based on details of Cohen's account that don't apply to my own (here I'm thinking primarily of Cohen's appeal to the notion of salience), some of Fantl and McGrath's critique would nonetheless be applicable to my current attempt to account for the wrongness of the relevant conjunctions, so a quick response is in order.

Fantl and McGrath most directly deal with me (specifically, with my 2005b, which has been incorporated into the present volume) here:

Keith DeRose (2005) claims that, when we are thinking about whether a third party subject may properly use a premise in practical reasoning (or to assert it), we 'should' use the standards for knowledge operative in the subject's own context. It is not clear to us that this is an appeal to salience. If it is, then it doesn't avoid the overridability problem. If it isn't, we don't understand why we should change our standards like this, unless KA is true. (2007: 577)

I hope that it is by now clear that, for me, this is just a matter of speakers using the standards appropriate to their conversational purposes. If you're talking about whether some subject 'knows' because you are wondering whether she might become a source by which you might also come to 'know' in a way that would help you in your practical situation, then your conversational interests in her will call for you to apply to her the standards relevant to your own practical situation. But if you're discussing whether she 'knows' in connection with evaluations of her own actions or potential actions (whether they're rational, prudent, moral, etc.), then of course your own conversational interests call on you to apply to her the standards relevant to her practical situation. That's what we should expect, given contextualism. To do otherwise is to use standards inappropriate to the very purpose for which you are evaluating whether the subject 'knows'. That seems sufficient to ground the wrongness of the relevant conjunctions—especially if the account takes the form of the third of my three options.

At (2007: 578–9), Fantl and McGrath charge contextualists who try to account for the wrongness of the relevant conjunctions in the way I pursue here, rather than by accepting and applying some unrelativized principle connecting knowledge with rational action, with operating in a way 'out of keeping with standard contextualist methodology', because, as they understand contextualists, we *are* quick to accept the closure principle for knowledge on grounds similar to those Fantl and McGrath are presenting for their knowledge-action principle. I address this below in n. 30.

more problematic than do (RP5) and (RP6), and this should be addressed. Hawthorne thinks the contextualist is committed to endorsing this sentence:

(H7) You should rely on propositions that you don't know to be true in your practical reasoning.[19]

Suppose that some contextualist, unlike me, was convinced that something like RKP is correct. Even so, it can seem that she will be committed to there being contexts relative to which (H7) is true. For, recall, whether it's appropriate for a subject to use a proposition in her practical reasoning depends entirely on that subject's own situation. Thus, the form of RKP that the contextualist will be able to endorse will be one that relativizes what subjects can use in their practical reasoning to what they count as knowing according to the epistemic standards appropriate to those subjects' own practical situations—something along these lines:

The Relativized Reason-Knowledge Principle (RKP-R): Where S's choice is p-dependent, it is appropriate for S to treat the proposition that p as a reason for acting iff S knows that p by the epistemic standards appropriate to S's practical situation.

And this connection between knowledge and practical reasoning, which is the only version of this type of connection the contextualist can avail herself of, can seem unable to rule out there being settings where (H7)—or at least something close to (H7)—is just the thing to say. For suppose a speaker is in a context governed by relatively high standards for knowledge. Still, there are agents in practical situations relative to which lower standards are appropriate who, consistent with RKP-R, are in a position to use propositions in their reasoning that they don't count as knowing by the higher standards that govern the speaker's context.

But Hawthorne is cheating in formulating the contextualist's apparent commitment as he does in (H7). Consider a basketball adviser who is constantly counseling coaches of basketball teams about the importance of playing someone tall at the center position. According to this adviser, that's one of the most important things to keep in mind while coaching. 'I really can't emphasize enough how important it is to have a tall player at center,' our adviser often says. Still, even this adviser admits there are some odd situations where the thing to do is to play someone who is not tall at center. Imagine

---

[19] Hawthorne (2004: 88). Hawthorne simply labels this sentence (7).

a critic of our adviser who, seizing on that admission, accuses our adviser of supporting the following recommendation:

(B7) You should use players who are not tall at the center position.

Surely our adviser has a legitimate complaint here! Not only is (B7) an unfair representation of our adviser's views because it is silent about her constant harping on the importance of tall centers, but, more to our current point, it also even inaccurately represents the one aspect of her views that it is narrowly focused on. Our adviser's admission doesn't commit her to (B7), but to something like:

(B7s) You should *sometimes* use players who are not tall at the center position.

For somewhat similar reasons, an accurate representation of the contextualist's apparent commitment is not to (H7), but to something like:

(H7s) It is *sometimes* appropriate for subjects to rely on propositions that they don't know to be true in their practical reasoning.

After all, it's not as if any contextualist would take a subject's belief falling short of certain epistemic standards to be a point *in favor of* the subject using the proposition that's the object of that belief in her practical reasoning, as (H7) can seem to suggest. Contextualists also have an apparent commitment to the following, which is quite similar to (H7s), and is an analogue of Hawthorne's (H6), adjusted to concern practical reasoning rather than assertion:

(R6) It often happens that people use propositions that they do not know to be true in their practical reasoning but they are not thereby subject to criticism.

The extent to which the contextualist is committed to even (H7s) and (R6) is unclear: If they sounded irredeemably implausible, the contextualist might be able to explain why they cannot be correct things to say in a way analogous to how, in the previous section, I explained why problematic assertions concerning knowledge and proper assertion were wrong. But the situation seems different here, for once the contextualist's even apparent commitments are presented fairly, they don't seem counter-intuitive: (H7s) and (R6) do not sound at all strange to me. In fact, they strike me as perfectly sensible statements.

Let's briefly consider a case that at least can seem to verify (H7s) and (R6). We can base our example on one that Hawthorne and Stanley consider as

the basis for an objection to RKP.[20] Suppose Ruth and Elizabeth are visiting a town that Ruth lived in many years ago, planning to look around a bit before going to Ruth's favorite restaurant in town, which they have verified is still in business. As they drive about, Ruth remembers more and more about where things are, and is now feeling fairly confident in her ability to find her way around. Ruth becomes hungry and starts to drive toward the restaurant, and as they come within a block of it, Ruth goes to the left at a fork in the road, feeling quite confident, though not completely certain, that the restaurant is that way. Imagine that the degree to which Ruth is now able to find her way around town is such that she is reasonable in feeling fairly confident that the restaurant is to the left. The following dialogue ensues:

Elizabeth: Hey, why did you go to the left at that fork? It looked like there were some interesting things to see down the road to the right.

Ruth: The restaurant was to the left, and I thought it was about time to eat. And look: Here it is.

E: We couldn't see the restaurant before we got to the fork. Did you know the restaurant was to the left when you went that way?

R: No, I didn't know it. But I was fairly confident. And it turns out I was right.

E: Wait! Let me get this straight. You didn't know that the restaurant was to the left, but you used the proposition that it was to the left in your practical reasoning?! You used that proposition that you didn't know to be true as a reason for acting?! Isn't that inappropriate?

R: No, I wouldn't say so. If I had been wrong, that would have just meant a bit more pleasant driving around before eating. No big deal. So I think it was perfectly reasonable for me to use that the restaurant was to the left in my practical reasoning in this case, even though I didn't quite know that to be so. I don't think I should be subject to criticism for that.

---

[20] The case that follows is based on one used by Hawthorne and Stanley in their (forthcoming), at their discussion of 'Objection 2' to their position in section III, 'Objections and Replies'. Both cases involve agents who don't know that the restaurant they are seeking is to the left, but who go to the left anyway. However, I have made my agent more confident (and reasonably so) that the restaurant is to the left: Hawthorne and Stanley describe their agent as having only a partial belief that the restaurant is to the left, and as merely thinking that 'it is somewhat more likely that the restaurant is to the left than to the right', while I'm about to describe my agent as reasonably feeling quite confident that the restaurant is to the left, even though she doesn't quite know that to be so. This change strengthens the challenge that such cases pose to (RKP), because it makes it more plausible to suppose that the agent is using the proposition that the restaurant is to the left, rather than some weaker proposition such as that the restaurant is *probably* to the left, as a reason for acting.

This dialogue takes an extremely unnatural turn toward its end, at Elizabeth's last contribution, where she moves the discussion to the fairly esoteric issue of whether it was appropriate for Ruth to use the proposition that the restaurant was to the left in her practical reasoning. Few non-philosophers would take a discussion like this in quite that direction. And, related to the unnaturalness of such discussions, I wouldn't claim to have any clear intuition about whether or not it is appropriate for Ruth to use the proposition that the restaurant was to the left in her practical reasoning, given that, as she admits, she didn't know that to be the case. What I do have a pretty clear intuition about is the reasonableness or rationality of Ruth's *action* of driving to the left: That was a perfectly rational thing for her to do in the situation, even if she didn't know that the restaurant was to the left. And I suppose it's then at least quite plausible to suppose that the practical reasoning an agent would typically use in a situation like Ruth's would be appropriate. What isn't entirely clear is whether that appropriate practical reasoning is best construed as relying on the proposition that the restaurant was to the left. In their defense of RKP, Hawthorne and Stanley handle such cases by construing the agents' practical reasoning such that other propositions that the subjects really do know, such as that the restaurant is *probably* to the left, constitute their reasons for acting.[21] So construing the agents' practical reasoning allows them to account for what are the fairly clear facts of such situations: that the agents are acting rationally and that the practical reasoning that typically leads to such actions is appropriate. So this case certainly does not provide a conclusive demonstration of the truth of (H7s) and (R6)—nor is it a conclusive refutation of RKP.

That said, Ruth's last line defending her practical reasoning, construed as relying on the proposition that the restaurant was to the left, even though she didn't know that proposition to be true, strikes me as perfectly sensible, which leads me to find (H7s) and (R6) quite believable. Such cases suggest this plausible general ruling: While there may be some situations (perhaps certain fairly high-stakes situations, for instance) where subjects should use only propositions they know to be true in their practical reasoning, there are also other cases, like Ruth's, where it is perfectly appropriate for subjects to so use propositions they do not know to be true. It is certainly at least

---

[21] Recall from the previous note the modification I have made to Hawthorne and Stanley's example. It's not certain that they would reply to this modified example in the same way that they responded to the one they did consider. But, if not, it isn't clear how they could respond to the modified example. The modification makes it even clearer that the agent is acting rationally in driving to the left, though she doesn't know that the restaurant is that way. Thus, it seems, protecting (RKP) from the example would seem to require finding *some* other proposition that Ruth does know, and that can be playing a key role in her practical reasoning in the place of the proposition that the restaurant is to the left.

very far from clear that even agents in low-stakes situations like Ruth's can *never* appropriately use propositions that they fall short of knowing as reasons for acting. I remain quite open to arguments that might show that in all cases like the one we have looked at here, it is only weaker propositions (e.g. concerning what is *likely* the case, rather than what is the case) that the subjects involved really do know that they appropriately can use in their practical reasoning.[22] But in the meantime, such cases and considerations give us good reason to be dubious of arguments that would take *as their starting point* that the likes of (H7s) and (R6) are unacceptable—and to be dubious of principles like RKP. Indeed, my suspicion is that Hawthorne himself would not traffic in arguments that take the unacceptability of the likes of (H7s) and (R6) as their starting points. Recall that, as he gave it, the argument in question starts from the implausibility of (H7), which he takes the contextualist to be committed to. And, indeed, (H7) does sound quite bad, at least to me. But as I've pointed out, (H7) is not an accurate representation of any even apparent commitment of contextualism: Contextualists instead have some apparent commitment to the likes of (H7s) and (R6), which can at least seem perfectly sensible.

The contextualist has little to fear from this part of Hawthorne's case against contextualism: The troubling sentence Hawthorne produces is not one the contextualist is committed to; once it is modified to accurately reflect the contextualist's (at least apparent) commitments, it is not troubling; and to the extent one *does* find it somewhat troubling, there seem to be good prospects for a contextualist account of why it would seem as implausible as it does.

9. Can the Contextualist Claim that Knowledge Is the Norm of Assertion?

To what extent does providing a contextualist account of the wrongness of troubling sentences like those treated in the previous two sections resolve Hawthorne's root worry that contextualism problematically breaks the ties that knowledge bears to proper assertion and to practical reasoning? That isn't entirely clear to me. Hawthorne could have been thinking of contextualism's endorsement of those sentences as particular manifestations of the underlying disease of problematically broken ties in such a way that dissolving

---

[22] See Hawthorne and Stanley (forthcoming) for an excellent discussion.

contextualism's apparent problem with those sentences wouldn't suffice to show that it doesn't suffer from the underlying disease. But in that case it isn't clear exactly what the underlying disease is supposed to be. As we saw in section 5, contextualism does allow for some important connections here. In what ways, then, does contextualism not allow for ties that seem to exist? That contextualism would seem to endorse some troubling sentences like those treated above seems to be at least the main reasonably specific complaint that Hawthorne registers here.[23]

But another worry along these lines that one might have about contextualism—and one that Hawthorne himself seems to have—is that there might be *principles* linking knowledge to assertion and to practical reasoning that contextualism will not allow for, at least in a full-blooded enough way, but which one might find very attractive. We'll discuss principles linking knowledge to proper assertion here, and ones linking knowledge to practical reasoning in the following section.

As I've already mentioned in section 6, as we saw in Chapter 3, and as Hawthorne realizes, contextualism not only allows for, but works very nicely with, this principle linking knowledge to warranted assertion:

> The Relativized Knowledge Account of Assertion (*KAA-R*): A speaker, S, is well enough positioned with respect to p to be able to properly assert that p if and only if S knows that p *according to the standards for knowledge that are in place as S makes her assertion.*

Yet Hawthorne has a point when he writes that 'the contextualist cannot quite claim that knowledge is the norm of assertion' (2004: 88). Note Hawthorne's use of 'quite': He doesn't seem ready to forthrightly declare that the contextualist simply can't claim that knowledge is the norm of assertion. I wish to be similarly 'hedgy' in my agreement with Hawthorne. But I admit in a comparative vein that the contextualist seems to have less of a right to that claim than an invariantist can have. According to the invariantist, there is only one knowledge relation—only one relation between subjects and propositions that is invoked by uses of the verb 'to know'. On classical invariantism, a subject is in that relation to a proposition at least roughly when she believes it, it is true, and she meets the single set of epistemic standards that govern all attributions of knowledge. So the classical invariantist can very straightforwardly

---

[23] For instance, at (2004: 87), this is what Hawthorne writes immediately after presenting (H6) (which, for Hawthorne, is simply numbered (6)), saying the contextualist is committed to it: 'Yet it sounds very odd. The link between knowledge and assertability has been severed.' This would seem to indicate that contextualism's commitment to the likes of (H6) at least largely *constitutes* the way in which Hawthorne thinks the view problematically severs 'the link between knowledge and assertability'.

tie proper assertion to knowledge by holding that a subject is in an adequate epistemic position to assert some proposition if and only if she stands in the one-and-only knowledge relation to that proposition—if and only if she (has a true belief and) meets the single set of standards for knowledge. But as we've seen, the daunting problem with this position is that the epistemic standards that govern proper assertion vary according to the speaker's context. The subject-sensitive invariantist like Hawthorne can seem to get the best of both worlds here, by allowing for varying standards for proper assertion, while at the same time holding that there is only one knowledge relation, which allows him to straightforwardly tie the standards for proper assertion to that single knowledge relation. He does this by understanding the knowledge relation in a different way: Instead of construing the single knowledge relation as one that a subject stands in with respect to a proposition iff she meets the particular epistemic standard that in fact provides the truth-condition for all our knowledge attributions, the subject-sensitive invariantist construes it as the relation a subject stands in with respect to a proposition iff she meets whatever epistemic standard is relevant to her current practical situation. That's a different—and, from certain angles of viewing it, a slightly more complicated—relation than that posited by the classical invariantist, but there is still, on SSI, only one knowledge relation, so the advocate of SSI can straightforwardly tie proper assertability to the issue of whether the subject stands in that one relation to the proposition in question.

By contrast, on contextualism, there are many different knowledge relations—however closely related they may be to one another. Sentences of the form 'S knows that p' express different relations, depending on the contexts in which they are asserted. Thus, the contextualist cannot straightforwardly tie a subject's being in a position to assert a proposition to her standing in the one-and-only knowledge relation. To the extent that the contextualist can in some way (kinda) make knowledge the epistemic norm for assertion, it will be by endorsing some principle like KAA-R, as I have done. But there are other tweakings of KAA that would also tie proper assertion to the various of the knowledge relations in various ways, even though they are not remotely as plausible as KAA-R. Below are just three such alternatives—one that we looked at briefly in Chapter 3, plus a very strong and a very weak account of assertability:

KAA-R2: A subject, S1, is well enough positioned with respect to p to make true the claim of a speaker, S2, that S1 is 'warranted in asserting that p' if and only if S1 knows that p according to the standards for knowledge that are in place in S2's context as S2 makes her claim.

KAA-RS: A speaker, S, is well enough positioned with respect to p to be able to properly assert that p if and only if S knows that p according to every epistemic standard that can ever govern assertions of 'S knows that p'.

KAA-RW: A speaker, S, is well enough positioned with respect to p to be able to properly assert that p if and only if S knows that p according to at least one of the epistemic standards that can ever govern assertions of 'S knows that p'.

As Hawthorne realizes, KAA-R is the most plausible way for the contextualist to somehow tie proper assertion to the various knowledge relations that the contextualist recognizes.[24] But that there are other possible ties that, though they are less plausible than KAA-R, have a claim equal to that of KAA-R's to having made knowledge the norm of assertion makes vivid the way in which advocates of KAA-R (like advocates of the other relativizations, if there were any) cannot simply and straightforwardly claim to have made knowledge the epistemic norm of assertion the way an invariantist can.

But that in itself is not much of an objection to contextualism or much of a comparative advantage to views like SSI over contextualism, unless some reason is given for thinking that a simple and straightforward tie is called for here. We have to look at the evidence for the knowledge account of assertion to determine if any of it favors the kind of straightforward and simple tie that invariantism makes possible over the contextualist's use of KAA-R. It is not sufficient for the objector to contextualism to cite evidence for some kind of connection between knowledge and proper assertion, infer that knowledge is the norm of assertion, and then fault contextualism for not being able to simply and straightforwardly make knowledge the norm of assertion, if the contextualist's proposed tie (KAA-R) can also account for the evidence ultimately being cited for the connection in the first place. The objector must

---

[24] Hawthorne writes: 'Let us consider the case of assertion from a contextualist perspective. If I have low standards and you have high standards, is it the property expressed by 'knowledge' in your mouth or the property expressed by 'knowledge' in my mouth that is the norm for your assertions? Plausibly, it is the former: If your assertions don't pass the standards for 'knowledge' in your mouth, then you shouldn't be making them. Correlatively, you should count as proper assertions of mine that can truly be labeled 'knowledge' by me even though they do not pass your high standards' (2004: 88).

It's perhaps worth clarifying the temporal relation between Hawthorne's (2004) and my (2002a), on which Chapter 3 of this book is based, and which relativizes KAA in the way specified by KAA-R (though I didn't use the abbreviation 'KAA-R' there, nor did I there discuss other possible relativizations) (2002a: 182). Though my (2002a) bears an earlier date than does Hawthorne's (2004), these two works should be considered contemporaneous (I don't claim any temporal priority for my work). My recollection is that we were both working on these publications at the same time, and we discussed these issues with one another while we were both working on the publications in question, discovering that we had some important ideas in common about the relation between knowledge and assertion, together with some important differences of opinion.

instead point to some particular evidence in the vicinity that the contextualist cannot account for—or cannot account for as well as the invariantist.

It seems to me that by wielding KAA-R, the contextualist is able to give an excellent account of the evidence typically cited for the knowledge account of assertion. Recalling from Chapter 3, that's particularly clear in the case of the evidence for KAA that I'm most moved by—namely by the knowledge account's ability to account for the clashes involved in the following types of conjunctions (all given in general schematic form, except for 4, which is an example of the type of conjunction, involving a 'simple "might"', that I mean to indicate):

1. P, but I don't know that p. Example: Fido barks, but I don't know that he does.
2. It's possible that not-p, but p. Example: It's possible that Fido doesn't bark, but he does bark.
3. It is possible that X V's, but it is not possible for X to V. Example: It is possible that Fido barks, but it is not possible for Fido to bark.
4. Fido barks, but he might not bark.
5. If p had been the case q would have been the case, but if p had been the case q might not have been the case. Example: If Fido had been there he would have barked, but he might not have barked if he had been there.

We have strong reason to think each of 1–5 is consistent.[25] Yet they all 'clash': They sound like contradictions. The hypothesis that a speaker represents herself as knowing a proposition by flat-out asserting it provides accounts of these seeming inconsistencies—most straightforwardly in the case of 1, but also, in conjunction with independently motivated accounts of various modal terms, with respect to 2–5.[26] These accounts of the apparent inconsistencies are very attractive because they support our sense that *some* inconsistency is responsible for the clash involved in asserting each of the types of conjunction, while, at the same time, happily removing that inconsistency from the realm of what's asserted: On all these accounts, the conjunction asserted is itself perfectly consistent, but in trying to assert it, one gets involved in a contradiction between one thing that one asserts and another thing that one merely represents as being the case. None of this explanatory power is lost when the hypothesis is construed along the lines indicated by KAA-R—that in asserting a proposition, a speaker represents herself as knowing that proposition by the epistemic standards that are governing her talk as she makes her assertion.

[25] See n. 21 of Chapter 3.
[26] See again n. 21 of Chapter 3 for references to where I work out these accounts.

Likewise, the contextualist's use of KAA-R seems to well account for the other good grounds that are typically cited in favor of the knowledge account of assertion.[27] But it's perhaps best not to explain why that's so with respect to all the good types of grounds that have been offered in support of KAA. It is for the objector to explain what grounds they are claiming can be better accounted for by the invariantist than by the contextualist who advocates KAA-R. What I am pointing out here is that absent such a case, the mere fact that the contextualist does not make knowledge the norm of assertion in as straightforward a way as can the invariantist does not constitute an objection to contextualism nor a relative advantage for the invariantist.

Or at least not much of an objection and not much of a relative advantage. Supposing that both the contextualist who utilizes KAA-R and an advocate of SSI who posits a more straightforward connection between knowledge and assertability can account for the data that support KAA, a slight other-things-being-equal preference might be assigned to the invariantist here in virtue of his somewhat simpler hypothesis. But such a slight advantage is easily overcome. That SSI is itself subject to strong objections of the type we looked at in sections 1–4 would do the trick here. But perhaps more directly relevant to our current concerns is the fact that the contextualist can actually better account for what ties there are connecting knowledge to proper assertion and to practical reasoning, for reasons we will see in section 12, below.

10. Principles Connecting Knowledge with Action

We have just seen how a contextualist can accommodate the grounds for KAA. Were there a case for a principle connecting knowledge with action that was as good as the one we have for KAA, a similar contextualist accommodation would be available there, too. However, there does not seem to be nearly as powerful a case for any such analogous principle connecting knowledge with action.

Recall (from section 6, above) Hawthorne and Stanley's principle concerning the connection between knowledge and action:

> *The Reason-Knowledge Principle (RKP)*: Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting iff you know that p.

---

[27] Here I'm thinking of the types of grounds given by Unger and by Williamson in the works cited in nn. 13 and 14 of Chapter 3.

For the reasons we saw in section 8, I am skeptical that anything like this principle is correct: Cases like the one we considered involving Ruth and Elizabeth driving to the restaurant argue against the principle and suggest the rival hypothesis that it's only in some situations that one should only take propositions one knows to be true as reasons for acting. I don't see how RKP gains a firm upper hand over this rival hypothesis. However, if there were good grounds for accepting something like RKP, then, depending on the exact nature of those grounds, the contextualist might well be able to accommodate them by accepting the relativized version of RKP we've already seen:

> *The Relativized Reason-Knowledge Principle (RKP-R)*: Where S's choice is p-dependent, it is appropriate for S to treat the proposition that p as a reason for acting iff S knows that p by the epistemic standards appropriate to S's practical situation.

Given the plausible expectation that speakers will utilize epistemic standards appropriate to their conversational purposes, and thus, when discussing whether or not subjects 'know' various propositions in connection with evaluations, predictions, or explanations of actions by those subjects, they will utilize standards appropriate to the practical situations faced by the subjects they are discussing, RKP-R would seem to hold a lot of promise for a contextualist account of the ties that exist between knowledge and practical reasoning in ways similar to those we have already explored (especially in section 7) by which the contextualist, utilizing KAA-R, can account for the data supporting connections between knowledge and proper assertion.

A big part of the problem with RKP is that it is a biconditional that is especially questionable in one of its two directions. The most troublesome cases for RKP are ones that, like our example of Ruth and Elizabeth, call into question whether knowing that p is *necessary* for the proper use of p in one's practical reasoning. More plausible than RKP, then, would be a principle that is like it, except that it claims only that knowledge that p is *sufficient* for being in position to use p in one's practical reasoning—or that in some way makes knowing a proposition sufficient for acting on it. We will call these 'one-way' knowledge-action principles. One such principle results from following Hawthorne and Stanley *very* closely and simply changing 'iff' to 'if' in their RKP, yielding:

> *1RKP-HS:* Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting if you know that p.

However, though I will pursue considerations that don't turn on matters of precise formulation while evaluating such principles, I think it is better to formulate such a principle as follows:

> *1RKP:* If an agent knows that p, then she is well enough positioned with respect to p to appropriately treat p as a reason for acting.

1RKP focuses solely on the matter of whether an agent is in a good enough epistemic position to treat p as a reason for acting, leaving open what other reasons there might be for it to be inappropriate for the agent to treat p as a reason for acting. (And it seems to be there may be other reasons beyond that the agent's choice is not 'p-dependent' in Hawthorne and Stanley's sense.) And what 1RKP says about this matter is that the agent's knowing that p is *sufficient* for S being in a position to treat p as a reason for acting, but it makes no claims about the necessity of knowledge here.

These one-way principles are considerably more plausible than two-way principles like RKP, because they refrain from endorsing what is by far the more problematic of RKP's directions. It is much harder to construct a convincingly correct dialogue in which a speaker admits that some subject knows something to be the case, but then, while that admission is still hanging in the air, goes on to complain that it was nonetheless inappropriate for the subject to take the proposition in question to be a reason for acting because she wasn't well enough positioned with respect to that proposition, without at least seeming to be taking back or canceling the admission that the subject knows. Such a complaint at least usually seems not to sit well with the concession that the subject knows—though we will soon consider examples in which such a complaint and such a concession don't sound awful together.

Thus, Jeremy Fantl and Matthew McGrath, who in recent work (2002, 2007) advocate a one-way principle similar to the ones we have been considering, but who seem to want no part of any two-way principles, are well motivated in their selectivity.[28] The one-way principle they defend and wield is the following:

> *KA*: S knows that p only if S is rational to act as if p.

---

[28] Though the cited papers by Fantl and McGrath are 'recent,' at least by my standards, their (2002) is to my knowledge the earliest defense of SSI. I hope my presentation here doesn't leave the impression that Fantl and McGrath were late-comers to the SSI party. As far as I can tell, they were first. My reason for beginning my exposition with Hawthorne, and largely focusing on him in this book, is that he is the source of the most pointed criticisms of contextualism from within the SSI camp. Hawthorne's challenging and prominent criticisms of contextualism make him the advocate of SSI (to the extent that he really is an advocate of SSI: see n. 12, above) that a contextualist such as myself has the most to answer to.

This is similar to 1RKP, but instead of focusing on which propositions agents can appropriately treat as reasons for acting, KA involves the (closely related) matter of which propositions are such that an agent can rationally act as if they are the case.[29] Though such one-way principles are considerably more defensible than their two-way counterparts, in (2002), Fantl and McGrath show how such one-way principles can nevertheless be quite effectively used in support of SSI.

But I also have doubts concerning whether something much like even these one-way principles can be correct. Let's look at the type of case that can challenge them. Jessica Brown proposes the following cases, both involving very high stakes:

SURGEON

A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetised on the operating table. The operation hasn't started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records.

AFFAIR

A husband is berating his friend for not telling him that his wife has been having an affair even though the friend has known of the affair for weeks.

Husband: Why didn't you say she was having an affair? You've known for weeks.

Friend: Ok, I admit I knew, but it wouldn't have been right for me to say anything before I was absolutely sure. I knew the damage it would cause to your marriage. (Brown 2008: 176–7)

These cases are attempts to do what I've said is difficult: to construct a convincingly correct dialogue in which a speaker combines an admission that a subject knows something to be the case with a complaint that it was nonetheless

---

[29] Fantl and McGrath (2007: 559). Fantl and McGrath give this gloss on their condition: 'To act as if p is to do whatever one is rational to do on the assumption of p. KA thus tells us that S knows that p only if S is rational to do whatever S is rational to do on the assumption of p.' In their (2002), Fantl and McGrath defend a slightly different principle formulated with the phrase 'preferring as if p' rather than 'acting as if p'; see Fantl and McGrath (2007: 559 n. 2) for the relation between these formulations.

inappropriate for the subject to act on that proposition because the subject wasn't well enough positioned with respect to it.

But are these convincingly correct? They sound fairly good to me, but these are like cases I've long been trying out on various people, and reactions to such cases seem to vary quite a bit. While the lines Brown puts into the mouths of the nurse and the friend certainly don't sound anything close to clearly wrong, they can also seem to fall short of being clearly correct, in part because there are replies available to these speakers that seem to be better, and seem to be better precisely because they avoid combining an admission that the subject knows something with a complaint about her acting on that proposition. In SURGEON, for instance, this reply seems available:

Nurse: Imagine what it would be like if she removed the wrong kidney! This is too important for her to be relying on her memory from this morning. She shouldn't operate before reviewing the patient's records one last time to be absolutely sure that she is removing the correct kidney.

It's interesting that this can seem such a good reply, since it sidesteps the question the student explicitly asked about whether the surgeon knows, and instead combines a negative judgement about the surgeon's epistemic position put in the different terms of her not being 'absolutely sure' with an evaluation of what course of action the surgeon should take. Similar replies are available in AFFAIR. In AFFAIR the husband explicitly states that the friend has known for weeks, and in SURGEON, the student explicitly asks in a pointed way whether the surgeon knows which kidney is diseased. If we modify the cases so that the questioner doesn't explicitly put things in terms of knowledge, replies which *deny* that the subject knows become available. Consider, for instance, this variation on SURGEON:

Student: I don't understand. Why is she looking at the patient's records again? Wasn't she in clinic with the patient just this morning?

Nurse: Imagine what it would be like if she removed the wrong kidney! This is too important to be relying on her memory from this morning. She shouldn't operate before checking the patient's records one last time so she can know that she is removing the correct kidney.

This reply that suggests the surgeon doesn't know before checking one last time, and a similar reply in the case of AFFAIR can compete with, and perhaps even edge out, in terms of apparent correctness, the replies that Brown considers.

On the other hand, as I've already noted, the replies Brown puts in her repliers' mouths also sound *fairly* plausible and natural, at least to me, leaving me in a state of some uncertainty about the prospects for one-way principles of the type we're considering.

Without trying to reach a definite final verdict on one-way principles,[30] I will present my contextualist take on the whole situation, which, while a bit wishy-washy in its current uncertain state, should be sufficient for the purpose of showing that the contextualist does not have much to fear here.

The important point here is the same as the one I made at the start of this section about alleged two-way principles: Insofar as one finds the grounds for these one-way principles to be convincing, then, depending on the exact nature of those grounds, the contextualist might well be able to accommodate these grounds by accepting relativized versions of those one-way principles, such as:

> *1RKP-R:* If an agent knows that p by the epistemic standards appropriate to her practical situation, then she is well enough positioned with respect to p to appropriately treat p as a reason for acting.

> *KA-R:* S knows that p by the epistemic standards appropriate to S's practical situation only if S is rational to act as if p.

Such relativized one-way principles, when combined with plausible expectations that speakers will utilize epistemic standards appropriate to their

---

[30] Several sections from a draft of this chapter, in which I pursued my doubts about one-way principles, evaluated the case for some kind of one-way knowledge/action principle in light of these doubts, and critically compared the grounds for such a principle with the grounds for some version of KAA, have been cut here largely because of length considerations. But I was also motivated by the fact that I couldn't reach a conclusive enough verdict: Including the material would add too much difficult reading to this chapter, all for the sake of too tentative a conclusion concerning this matter.

However, one matter dealt with in the cut material should be addressed here, even if only briefly. At (2007: 578–9), Fantl and McGrath charge contextualists who do not accept their unrelativized KA with operating in a way 'out of keeping with standard contextualist methodology', because, as they read us, we *are* quick to accept the closure principle for knowledge on grounds similar to those Fantl and McGrath are presenting for KA. And I know that I, for one, am widely considered to be a fervent proponent of closure. However, in a way that is relevant to Fantl and McGrath's charge, my own position on the issue of closure is in fact a very tempered one. (In DeRose 1995, the paper of mine on which my pro-closure reputation seems to be largely based, I mention closure principles only in one footnote (DeRose 1995: 31–2 n. 33), and am there rather tepid in my response to such principles.) While I do reject certain 'abominable conjunctions' that run afoul of closure principles (DeRose 1995: 28), I do so because of their abominableness, and not because I am antecedently committed to some closure principle that they violate. And though I do suspect that there is *some* correct principle in the vicinity of closure principles that have been put forward, for reasons indicated at DeRose (1999c: 23 n. 14), I have yet to see a formulation of a closure principle that I can endorse. So, without pursuing all the details relevant to the comparison, let me just state that I think my cautious approach to knowledge/action principles is very much of a piece with my approach to closure principles.

conversational purposes in ways like those we explored in section 7, seem to hold much promise for accounting for a lot of the data that might be marshaled in favor of their non-relativized counterparts.[31] In particular, it seems quite promising so to account for the consideration that, three paragraphs above, I raised in favor of one-way knowledge-action principles: that it often seems somehow wrong, or at least very uncomfortable, for a speaker to combine a claim that a subject 'knows' a proposition to be the case with a complaint that that subject shouldn't have acted on that proposition (or taken that proposition as a reason for acting) because she wasn't well enough positioned with respect to the proposition to do so.

What do we say about examples like Brown's, then? To the extent that counter-examples like Brown's are convincing, one can maintain the one-way principles, but ease up on the accompanying pragmatics that one employs. Actually, the best course in that case would seem to be to revert to a relativized version of some two-way principle—like RKP-R, which, recall, is just like 1RKP-R, except for having 'iff' where 1RKP has only 'if'—but ease up on the accompanying pragmatics one appeals to in a somewhat uneven way. We still maintain that speakers should in general utilize epistemic standards that are appropriate to their conversational purposes, and this means that when they are discussing whether subjects 'know' propositions in connection with evaluations of those subjects' actions or potential actions, they should in general utilize the standards appropriate to the subjects' practical situations—the epistemic standards that represent how well positioned the subjects should be with respect to propositions before they are rational to act on those propositions. However, we can also recognize that a speaker may have reasons for utilizing a standard for 'know(s)' that differs from the standard appropriate to the subject's practical situation. Especially relevant to our current concerns is this potential reason: that the epistemic standards relevant to the subject's practical situation differ too much from the normal range of standards that usually govern 'know(s)', or from the standards one has some other reason to keep in place for 'know(s)'. In such a case, it can be permissible for a speaker to let the epistemic standards governing her use of 'know(s)' to diverge from those appropriate to the

---

[31] In the contextualist's account, the (relativized) principles themselves are not very substantial. If by 'the epistemic standards appropriate to the subject's situation' we mean the standards that must be met before the subject can properly use the proposition in question as a reason to act, then these principles become quite empty, though they do embody the assumption that what it takes to count as 'knowing' something, according to varying standards, correlates with varying standards that it takes to be in an epistemic position to properly act on a proposition. The pragmatics the contextualist employs is the real workhorse of the contextualist's account.

practical situation of the subject, so long as it is somehow clear that that is what is going on. So, if a subject's practical situation is such that it's rational for her to act on propositions that meet even very low standards, a speaker may opt to use 'know(s)' according to standards higher than those very low ones appropriate to the subject's situation, and talk about the subject as not 'knowing' propositions that she is nonetheless rational to act on: 'She doesn't know that p is true, but don't sweat it: Not much is riding on this, and it's fine for her to act on the proposition that p even though she doesn't know it to be the case.' That is what's happening in the last line of our dialogue in section 8, where Ruth defends her acting on her belief that the restaurant was to the left, even though she didn't know that to be the case. On the other hand, when the subject's practical situation makes very high standards appropriate, speakers may perhaps allow standards more moderate than those very high ones to govern their use of 'know(s)', and to describe subjects as 'knowing' propositions that the subjects cannot rationally act on, because their position with respect to those propositions isn't good enough. That is what's happening in the last lines of Brown's dialogues, insofar as those last lines are correct (as I myself am inclined to judge them). The relative difficulty in coming up with clearly correct cases of this type shows that in cases where subjects are being discussed in connection with evaluations of those subjects' own courses of action, it is harder to allow the standards governing one's use of 'know(s)' to go below the standards appropriate to the practical situation of the subject one is discussing than it is to let the standards rise above that level. Of course, I cannot claim to be able to anticipate all the data that might plausibly be cited and arguments that might be made in favor of either two-way or one-way knowledge-action principles; nor to confidently predict what version of such a principle, if any, will prove correct; nor whether any considerations that will be brought to bear on the matter will favor a non-relativized over a relativized version of the best principle. But from my admittedly limited current perspective, I don't see much reason to be pessimistic about the prospects for contextualism here. At the very least, I haven't yet seen a convincing objection to contextualism that has emerged from the type of considerations we have been considering in this chapter.

11. Multi-tasking and the Case of the Walking Talker

For the previous five sections (6–10), I have been playing defense. It's almost time to go on offense. It turns out that contextualism has a very important

advantage over SSI in accounting for uses of 'knows' made in connection with evaluating, predicting, and explaining actions, or so I will argue. I will make that case in the following section. To prepare for that case, we will in this section discuss an example of 'multi-tasking'.

So far, in discussing various cases and principles, I've been rather blithely writing of 'the standards appropriate to the practical situation' of various subjects. Things get interesting when we start to consider cases in which single subjects face multiple practical 'situations', where different epistemic standards are appropriate to evaluating the various potential actions and assertions in the various practical situations that these subjects face.

So, consider the case of Jane, the walking talker. This case is tailored to draw in advocates of SSI who tie knowledge very tightly both to proper assertion and to practical reasoning, and who will make both of the following judgements. First, they will hold that a young, healthy assistant professor like Jane, who is just sitting in her office and talking on the phone with a friend who is wondering out of fairly idle curiosity where Jane will be next year, is in a position to assert that, and also knows that, she will be back at work at her home institution in a year's time, and won't have to add to her assertion of that fact any proviso like, '…supposing, of course, that I'm still alive a year from now'. But, second, they will judge that when Jane faces a different practical situation where, say, she's deciding whether to take the five-minute walk to her college's benefits office to sign up for special (free) life insurance to cover her for the year that she's away, it can be rational for her to take the walk and obtain the insurance, and, tightly connected with that, that Jane will not know that she will survive the year, nor that she will be back at work at her home institution in a year's time, when she faces that kind of situation. Some advocates of SSI may not agree with these particular judgements, but they will all hold that different epistemic standards apply to subjects depending on the practical situations they face, so analogous combined cases should be available to target SSI advocates who don't subscribe to the particular judgements in play here.

Our case of the walking talker, of course, combines those two situations. Suppose that, as before, Jane is rationally taking her walk to the benefits office. She doesn't particularly need the exercise, and, were it not for the insurance she's getting, she'd rather use the time the task takes to do something else, but the special insurance is a good enough deal to make it rational for her to take the time needed to sign up for it. Now, as she's walking (and, we'll suppose, she continues her very rational walk throughout this episode), her cell phone rings, and, just as happened when we were supposing Jane was sitting in her office, it's her friend calling, just idly wondering where Jane will

be and what she will be doing in a year's time. It seems that, as before, Jane is in a position to assert to her friend that she'll be back at work at her home institution, and will not have to add any morbid proviso to her assertion for it to be warranted.

According to SSI, in this combined case in which she is both walking and talking, as in all cases, Jane still either *simply knows* that she will survive the year and will be back at her work in a year's time or she *simply doesn't know* those facts—or, perhaps, depending on their general approach to some issues in philosophical logic, advocates of SSI might opt for a third option in such cases of 'multi-tasking': that it's *simply indeterminate* whether Jane knows (or something along those lines).³²

But which is it? This question presents something of an embarrassment for an advocate of SSI who wants to tie knowledge to assertion and to action as tightly as Hawthorne wants to, for each of the answers threatens one of those tight ties and implicates one in a sentence that will seem about as strange as the sentences Hawthorne charges contextualism with endorsing. Supposing, as is extremely plausible, that Jane is as rational in taking the time to sign up for the insurance and is as warranted in making her assertion in this combined case as she was with respect to each of these actions in their separate cases, taking the first option will lead one to say that in the combined case it's rational for Jane to take the time to sign up for special life insurance that will cover her only in the case of her death and only for the coming year, even though she *knows* (and, in fact, *simply knows*) that she will survive that year. That sounds little if any better here than when it is said of single-tasking Jane who receives no phone call as she walks to the benefits office. Taking the second option instead leads one to claim that Jane is warranted in flat-out asserting to her friend some facts that Jane does not know (and, in fact, that she *simply does not know*) to be the case. That sounds little if any better than when it is said of single-tasking Jane who is just sitting in her office as she talks to her friend. The third option may soften these blows a bit, but it doesn't seem to help the plausibility of

---

³² On our semi-technical use of it, 'simply' is not used as an emphatic endorsement of bivalence (though it *could* be so used in the likes of 'someone either simply knows or simply doesn't know', it seems to me). Rather, it signifies a form of speaker-independence in the truth-value of the sentence in question. On our use, to say that a subject 'simply knows' is to say that the subject is such that any speaker, no matter their context, will be asserting a truth if they say the subject 'knows', and will be asserting something false if they instead say that the subject does not 'know'. Analogously, to say that it's 'simply indeterminate' whether a subject knows is to say something like that the subject is such that any speaker, no matter their context, will be saying something that's neither true nor false if they say either that the subject 'knows' or that she doesn't 'know'.

one's position all that much to instead combine the plausible evaluations of Jane's action (that it's rational) and assertion (that it's warranted) with the claim that it's simply indeterminate whether or not Jane knows the relevant facts upon the knowledge of which that rationality and warrantedness seems to turn.

Though, as I said, our combined case of the walking talker may be something of an embarrassment to versions of SSI that tie knowledge in some very tight ways to rational action and to proper assertion, the example does not in the final analysis provide a very strong objection to such views, because, when one stands behind both of the intuitively very plausible judgements that Jane's action is rational and her assertion is warranted in the combined case, then *intuitively and independent of any particular theoretical commitments one might have*, it's difficult to answer a pointed question about whether a subject like Jane knows that it is made in reference to two such actions of hers (the assertion and the walk to the benefits office) that she is very saliently engaged in and that seem to call for different standards. Suppose that someone is listening to you telling Jane's story, and is disturbed by your endorsement of both Jane's action and her assertion, and so interrupts with this pointed challenge: 'Wait a minute! You seem to think it's perfectly fine for Jane to flat-out assert to her friend that she'll be back at her school in a year, but that it's also perfectly rational for her to take the time to sign up for the special insurance that will cover her only for that year. So now, tell me, does she know that she'll survive the year and will be coming back to her school, or not? Which is it?' There are various ways of responding to such questions that are more or less intuitively plausible, but I have found that what answers people find more or less plausible varies quite a bit among individuals, and there's no clearly great way to answer such questions. Therefore, that adherents to a certain theoretical position may have a hard time dealing with such a case is not a strong objection to their position. What seems to *me* the best way to respond is to say something that involves explicit relativization of Jane's knowledge to the different standards involved, as in: 'Well, Jane does know that she'll survive the year and be back well enough to assert to her friend that she will be back, but she doesn't know that she'll survive with enough certainty to make it irrational for her to take the trouble to sign up for the special life insurance.' But this response, in addition to being unsatisfying to some who would press the pointed question (and who might still insist on being told whether Jane knows or not, and might not agree that they've already been told everything there is to tell here, as the contextualist is likely to inform them), seems available whether SSI or contextualism is accepted.

12. Contextualism's Advantage over SSI in Accounting for Uses of 'Know(s)' Made in Connection with Evaluations, Predictions, and Explanations of Actions

It may be tempting to dismiss cases of 'multi-tasking' as rare cases about which it's just hard to know what to say. But as we are about to see, (a) in a way that's relevant to our assessment of contextualism and of SSI, it is *extremely* common for agents to be multi-taskers, and (b) there are some clear facts about our use of 'know(s)' in evaluating (and predicting and explaining) the actions of such multi-taskers that contextualism accounts for, but SSI cannot.

On point (a), when we look not just at what courses of action agents are actually engaged in, but at what potential or possible courses of action they might face (or might have faced) in connection with which some speakers might take an interest in discussing them, then, of course, we're all constantly subject to massive multi-tasking: there are many possible practical situations relative to which some speakers might conceivably talk about us. And since speakers don't limit their talk of whether subjects actually 'know' things to discussions of those subjects' actual practical situations, but feel free to cite facts about whether subjects do or do not actually 'know' things to be the case in connection with merely possible practical situations that, for all the speakers know, the subjects might face (and indeed, even for situations the speakers do know the subjects won't face), the way in which massive multi-tasking is the reverse of uncommon is the very sense relevant to our current task of evaluating theories, since these theories apply to our talk of subjects' (actual) 'knowledge' in connection with both possible and actual practical situations that subjects might face.

The above points are perhaps most vividly illustrated by a very sparsely and abstractly described example. So, consider Judith, and some true proposition p that Judith believes and relative to which she is in a moderately strong epistemic position. Suppose there are many possible practical situations that, so far as her friends know, Judith might soon face, relative to the evaluation of which it may be relevant whether Judith actually 'knows' that p. Some of these possible situations are very high-stakes scenarios, relative to which very demanding epistemic standards seem relevant. Others are low-stakes situations, relative to which lower epistemic standards seem relevant. Which of these practical situations does Judith actually face? We may suppose that right now, as her friends realize, Judith is sound asleep. To one aware of the relevant facts about Judith, she doesn't seem to be facing any of these practical

situations in particular right now, and it's just unclear which, if any, of them she may be facing in the coming day or days. Since it isn't always clear just how to construe what a given subject's 'practical situation' is (how much it encompasses) and what epistemic standards are appropriate to it, what will be important to our argument as we imagine how various of Judith's friends in various conversations are likely to talk about Judith is that they all have the same information about what Judith's actual practical situation is—though we will imagine they have very different conversational interests in Judith. It will be fairly clear that, despite having the same information about what Judith's practical situation actually is, her friends will talk very differently about whether she 'knows' that p, in ways that we should find very surprising on the assumption of SSI.

What epistemic standards will Judith's friends employ in talking about what Judith 'knows' in connection with the various practical situations she might face? Moving now to point (b), the answer to this question is one that strongly favors contextualism—namely, that what standards they will apply depends on just what potential practical situation they happen to be discussing. For the advocate of SSI, Judith either simply knows or simply doesn't know (or perhaps they may hold that it may be simply indeterminate whether she knows): What standards apply to a subject is set by that subject's situation, and different speakers who have the same information about what the subject's actual situation is will not apply different standards to the subject based on differences in the speaker's own conversational contexts. But speakers clearly do display just the kind of freedom SSI would deny them: Depending on just what potential courses of action speakers are discussing, they will apply different standards to Judith, even where they have the same information about what Judith's actual practical situation is.

More fully described examples could be used to drive this point home, but I won't take the space to construct or discuss particular cases here, except for briefly, in a note,[33] because, first, the reader should be able to construct

---

[33] For example, we could use a variation of our Thelma/Louise/Lena stories in which Thelma and Louise have the same information, both about matters concerning how well positioned Lena is with respect to the proposition that John was at work, and about matters relevant to what practical situation Lena is actually facing, but in which it is utterly natural and seems utterly right for Thelma, in LOW (at the tavern), to say that Lena 'knows' that John was at work in support of a judgement that Lena would have stopped by to pick up her winnings on the bet if she had made a bet, while it is utterly natural and seems completely correct for Louise, in HIGH, to say Lena 'doesn't know' that same fact in support of a judgement that Lena wouldn't have been in a position to tell the police that John was at work. The trick here is to neutralize the difference that Louise is aware of the fact that the police might be interested in talking to Lena, while Thelma is not. But it's easy to tell the story in a way that gives them the exact same information about this matter, as I did above in section 2, and this

the cases for herself by now. And, second, this would all be in the service of making a point that should be clear upon reflection, anyway: that when discussing what a particular subject does and does not actually 'know' in connection with judgements being made about potential courses of action the subject might take or might have taken, different groups of speakers will apply different epistemic standards to the subject depending on what possible practical situations the speakers happen to be discussing, even where the different groups of speakers have all the same information pertaining to what practical situation the subject actually faces. In such cases, speakers clearly display the freedom that contextualism grants them, but that SSI would deny them: We are not bound to apply to subjects the standards appropriate to the practical situations those subjects actually face—to in that way talk of them as if they either simply know or simply don't know, as SSI would have it. Rather, we are quite free to apply different standards to the subjects, even with no change in our information about their actual epistemic or practical situations, depending on what conversational interests we as speakers take in those subjects—where, in these cases, the different interests consist in different potential courses of actions in relation to which we choose to discuss the subjects.

As I admitted in the previous section, when, all at the same time, a group of speakers discuss whether a subject 'knows' something to be the case in connection with two or more very different practical situations that the subject is actually engaged in and that call for quite different epistemic standards, it isn't completely clear what should be said (though my suggestion is to engage in some explicit relativization). It is therefore not much of a mark against a theory that it can't smoothly handle such situations. And that is not only true of situations like we considered in the previous section, where the subject is actually and saliently engaged in two such courses of action, but is also true if a subject is being discussed simultaneously in connection with two merely possible practical situations the subject might face, or in connection with one actual and one merely possible practical situation, where the evaluations of those practical situations seem to call for different epistemic standards. Again, it isn't clear what to say about such cases.[34] But what we've been considering

does not in the slightest change what it seems natural and right for the two speakers to say, so long as our two speakers' conversational interests differ, with Thelma discussing what Lena 'knows' only in connection with the low-stakes practical scenario, and Louise discussing the issue only in connection with high-stakes concerns. Variations on our old friends the Bank Cases could also be used here.

[34] For what it's worth, my own suggestion as to what's best to say in such cases is what I suggested for the cases in which the subject is actually engaged in two separate practical situations: explicit relativization, as in 'Well, he knows by the standards relevant to this [potential] task, but not by the standards relevant to that one.'

here in this section are situations in which an agent is being discussed in a given conversation in connection with just one practical situation they might face, where one set of epistemic standards seem relevant to the evaluations being made of the actions or potential actions of the agent in that potential practical situation. Here, it is pretty clear what speakers often can, do, and should say: They will apply to the subject the epistemic standards relevant to the particular practical matter they happen to be discussing. However, to the chagrin of advocates of SSI, it turns out that other speakers who are discussing the actual epistemic standing of the same agent in connection with another possible practical situation she might have faced will just as naturally apply to that subject the different epistemic standards relevant to the evaluations they happen to be making about her. The standards that speakers should and do apply to subjects depend on the contexts of the speakers in a way that SSI does not allow for, but that contextualism handles nicely.

13. The Need for the Flexibility Contextualism Posits

One natural way to view the difference between SSI and contextualism is that, while both views allow that different epistemic standards govern knowledge-attributing sentences in different situations, SSI rules that the subject's context sets the standards, while, as Hawthorne writes, according to contextualism, 'It is always the ascriber's standards that call the semantic shots, so to speak' (2004: 59). Thinking of the matter in this way, it can seem that the uses of 'know(s)' that we've been focused on since section 5—uses in which whether or not a subject 'knows' something is being discussed in connection with an evaluation, prediction, or explanation of an action or potential action of that subject—give us strong reason to favor SSI over contextualism. For in such cases, speakers naturally use 'know(s)' according to the standards that are appropriate to the practical situation that the subject they are talking about faces.

However, as we saw in sections 2–3, there are other contexts in which different conversational purposes are in play, and in which speakers will apply to subjects standards that are appropriate to the practical interests of the speakers, with no regard for what standards are appropriate to the practical situation being faced by the subject.

As remarked in section 4, all these considerations can appear to leave us with a big, ugly tie in the contest between the two views: each handles certain uses of 'know(s)' nicely, while fumbling other uses. At this point, some who can easily be tempted by such thoughts might develop the suspicion that there are

somehow two different senses of 'know(s)', one used for evaluating subjects as potential sources of information (and about which contextualism might be correct), and one used in connection with evaluating subjects as practical agents (about which SSI may be right), or something like that. Or perhaps some others are tempted at this point to think that, though 'know(s)' doesn't have two such senses, its meaning is incoherent because it somehow 'tries' to do two very different jobs.

Fortunately, the big ugly tie is broken by the realization that, on contextualism, speakers' own conversational interests can lead them to apply to subjects the standards that are appropriate to the practical situations faced by those subjects. (Or, more precisely, this realization shows that the two views were never really tied in the first place.)

A more enlightening way to view the difference between the two views emerges: According to SSI, a subject's practical situation determines the epistemic standard that applies to her, and speakers are constrained to apply that standard to the subject, while on contextualism, speakers can apply different standards to a subject, depending on the conversational interest the speakers take in the subject. When we've turned to the conversational phenomena, what we've seen in this chapter is that speakers seem to display the freedom that contextualism grants them and that SSI would deny them. For, first, again, while we *can* apply to subjects the epistemic standards appropriate to the practical situations those subjects face, and we will naturally do so when we take certain conversational interests in them, other clearly correct uses of 'know(s)' show that we can also take other interests in subjects, and apply epistemic standards to them that SSI would forbid us to apply. And second (and this may be important to any that might still harbor a suspicion that 'know(s)' has both a contextualist and an SSI sense), when we try to account for uses of 'know(s)' that are made in connection with subjects' own actions or potential actions, it turns out, for reasons we saw in section 12, that even here we need the freedom that contextualism allows but that SSI denies us.

References

Annis, David 1978. 'A Contextualist Theory of Epistemic Justification', *American Philosophical Quarterly* 15: 213–19.

Austin, J. L. 1979a. 'Other Minds', in Austin, *Philosophical Papers*, 3rd edn., Oxford: Oxford University Press, 76–116.

—— 1979b. 'A Plea for Excuses', in Austin, *Philosophical Papers*, 3rd edn., Oxford: Oxford University Press, 175–204.

Bach, Kent 1999. 'The Myth of Conventional Implicature', *Linguistics and Philosophy* 22: 327–66.

Black, Max 1952. 'Saying and Disbelieving', *Analysis* 13: 25–33.

Blackson, Thomas A. 2004. 'An Invalid Argument for Contextualism', *Philosophy and Phenomenological Research* 68: 344–5.

Boër, Stephen, and Lycan, William 1975. 'Knowing Who', *Philosophical Studies* 28: 299–344.

—— —— 1986. *Knowing Who*, Cambridge, Mass.: MIT Press.

Braun, David 2006. 'Now You Know Who Hong Oak Yun Is', *Philosophical Issues* 16: 24–42.

Brown, Jessica 2005. 'Adapt or Die: The Death of Invariantism?', *Philosophical Quarterly* 55: 263–85.

—— 2008. 'Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning', *Noûs* 42: 167–89.

Brueckner, Anthony 1994. 'The Shifting Content of Knowledge Attributions', *Philosophy and Phenomenological Research* 54: 123–6.

Cappelen, Herman, and Lepore, Ernie 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*, Malden, Mass.: Blackwell Publishing.

—— —— 2006. 'Replies [to Kent Bach; John Hawthorne; Kepa Korta and John Perry, and Robert J. Stainton]', *Philosophy and Phenomenological Research* 73: 469–92.

Christensen, David 2007. 'Epistemology of Disagreement: The Good News', *Philosophical Review* 116: 187–217.

Cohen, Stewart 1986. 'Knowledge and Context', *Journal of Philosophy* 83: 574–83.

—— 1987. 'Knowledge, Context, and Social Standards', *Synthese* 73: 3–26.

—— 1999. 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* 13: 57–89.

—— 2001. 'Contextualism Defended', *Philosophical Studies* 103: 87–98.

—— 2004. 'Contextualism and Unhappy-Face Solutions: Reply to Schiffer', *Philosophical Studies* 119: 185–97.

Crimmins, Mark, and Perry, John 1989. 'The Prince and the Phone Booth: Reporting Puzzling Beliefs', *Journal of Philosophy* 86: 685–711.

DeRose, Keith 1990. 'Knowledge, Epistemic Possibility, and Scepticism', Ph.D. diss., University of California, Los Angeles; University Microfilms International.

—— 1991. 'Epistemic Possibilities', *Philosophical Review* 100: 581–605.

—— 1992. 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* 52: 913–29.

—— 1994. 'Lewis on "Might" and "Would" Counterfactual Conditionals', *Canadian Journal of Philosophy* 24: 413–18.

—— 1995. 'Solving the Skeptical Problem', *Philosophical Review* 104: 1–52.

—— 1996a. 'Knowledge Assertion, and Lotteries', *Australasian Journal of Philosophy* 74: 568–80.

—— 1996b. 'Contextualism', in D. M. Borchert (ed.), *Encyclopedia of Philosophy: Supplement*, New York: Simon & Schuster Macmillan, 111–13.

—— 1996c. 'Relevant Alternatives and the Content of Knowledge Attributions', *Philosophy and Phenomenological Research* 56: 193–7.

—— 1998. 'Simple *Might*s, Indicative Possibilities, and the Open Future', *The Philosophical Quarterly* 48: 67–82.

—— 1999a. 'Contextualism: An Explanation and Defense', in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology*, Malden, Mass.: Blackwell Publishers, 187–205.

—— 1999b. 'Can It Be That It Would Have Been Even Though It Might Not Have Been?', *Philosophical Perspectives* 13: 385–413.

—— 1999c. 'Responding to Skepticism', in K. DeRose and T. A. Warfield (eds.), *Skepticism: A Contemporary Reader*, Oxford: Oxford University Press, 1–24.

—— 2000. 'Now You Know It, Now You Don't', in *Proceedings of the Twentieth World Congress of Philosophy*, v: *Epistemology*, Bowling Green, Oh.: Philosophy Documentation Center, 91–106.

—— 2002a. 'Assertion, Knowledge, and Context', *Philosophical Review* 111: 167–203.

—— 2002b. Review of Williamson (2000), *British Journal for the Philosophy of Science* 53: 573–7.

—— 2004a. 'The Problem with Subject-Sensitive Invariantism', *Philosophy and Phenomenological Research* 68: 346–50.

—— 2004b. 'Sosa, Safety, Sensitivity, and Skeptical Hypotheses', in J. Greco (ed.), *Ernest Sosa and his Critics*, Oxford: Blackwell Publishing, 22–41.

—— 2005a. 'Direct Warrant Realism', in A. Dole and A. Chignell (eds.), *God and the Ethics of Belief: New Essays in Philosophy of Religion*, Cambridge: Cambridge University Press, 150–72.

—— 2005b. 'The Ordinary Language Basis for Contextualism and the New Invariantism', *Philosophical Quarterly* 55: 172–98.

—— 2007. Review of Stanley (2005), *Mind* 116: 486–9.

—— 2008. 'Gradable Adjectives: A Defence of Pluralism', *Australasian Journal of Philosophy* 86: 141–60.

—— and Grandy, Richard E. 1999. 'Conditional Assertions and "Biscuit" Conditionals', *Noûs* 33: 405–20.

Donnellan, Keith 1966. 'Reference and Definite Descriptions', *Philosophical Review* 75: 281–304.

—— 1968. 'Putting Humpty Dumpty Together Again', *Philosophical Review* 77: 203–15.

—— 1978. 'Speaker Reference, Descriptions, and Anaphora', in *Syntax and Semantics*, vol. ix, Academic Press. Page references to the reprint in P. French, T. Uehling, and H. Wettstein (eds.), *Contemporary Perspectives in the Philosophy of Language* (Minneapolis: University of Minnesota Press, 1979), 28–44.

Douven, Igor 2006. 'Assertion, Knowledge, and Rational Credibility', *Philosophical Review* 115: 449–85.

Dretske, Fred 1970. 'Epistemic Operators', *Journal of Philosophy* 67: 1007–23.

—— 2005a. 'The Case against Closure', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*, Oxford: Blackwell Publishing, 13–26.

—— 2005b. 'Reply to Hawthorne', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*, Oxford: Blackwell Publishing, 43–6.

Dudman, V. H. 1992. 'Probability and Assertion', *Analysis* 52: 204–11.

Elga, Adam 2007. 'Reflection and Disagreement', *Noûs* 41: 478–502.

Fantl, Jeremy, and McGrath, Matthew 2002. 'Evidence, Pragmatics, and Justification', *Philosophical Review* 111: 67–94.

—— —— 2007. 'On Pragmatic Encroachment in Epistemology', *Philosophy and Phenomenological Research* 75: 558–9.

Feldman, Richard 1999. 'Contextualism and Skepticism', *Philosophical Perspectives* 13: 93–115.

—— 2006. 'Epistemological Puzzles about Disagreement', in S. Hetherington (ed.), *Epistemology Futures*, New York: Oxford University Press: 216–36.

—— and Warfield, Ted (eds.) forthcoming. *Disagreement*, Oxford: Oxford University Press.

Gettier, Edmund L. 1963. 'Is Justified True Belief Knowledge?', *Analysis* 23: 121–3.

Goldman, Alvin 1976. 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73: 771–91.

—— 1992. *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, Cambridge, Mass.: MIT Press.

Grice, H. P. 1961. 'The Causal Theory of Perception', *Proceedings of the Aristotelian Society* Supplementary Vol. 35: 121–52.

—— 1989. *Studies in the Way of Words*, Cambridge, Mass.: Harvard University Press.

Halliday, Daniel 2007. 'Contextualism, Comparatives, and Gradability', *Philosophical Studies* 132: 381–93.

Hambourger, Robert 1987. 'Justified Assertion and the Relativity of Knowledge', *Philosophical Studies* 51: 241–69.

Hawthorne, John 2000. 'Implicit Belief and A Priori Knowledge', *Southern Journal of Philosophy* 38, Spindel Conference Supplement: 191–210.

—— 2004. *Knowledge and Lotteries*, Oxford: Oxford University Press.

—— 2005. 'The Case for Closure', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*, Oxford: Blackwell Publishing, 26–43.

—— 2006. 'Testing for Context-Dependence', *Philosophy and Phenomenological Research* 73: 443–50.

——— and Stanley, Jason forthcoming. 'Knowledge and Action', *Journal of Philosophy*.

Hazlett, Allan forthcoming. 'The Myth of Factive Verbs', *Philosophy and Phenomenological Research*.

Heck, Richard G. 2006. 'MacFarlane on Relative Truth', *Philosophical Issues* 16: 88–100.

Holton, Richard 1997. 'Some Telling Examples: A Reply to Tsohatzidis', *Journal of Pragmatics* 28: 625–8.

Horn, Laurence R. 1989. *A Natural History of Negation*, Chicago: University of Chicago Press.

Jackson, Frank 1979. 'On Assertion and Indicative Conditionals', *Philosophical Review* 88: 565–89.

Johnson, E., Hershey, J., Meszaros, J., and Kunreuther, H. 2000. 'Framing, Probability Distortions, and Insurance Decisions', in D. Kahneman and A. Tversky (ed.), *Choices, Values, and Frames*, Cambridge: Cambridge University Press.

Kaplan, David 1989. 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan*, Oxford: Oxford University Press, 481–563.

Kelly, Thomas 2005. 'The Epistemic Significance of Disagreement', *Oxford Studies in Epistemology* 1: 167–96.

Kripke, Saul 1977. 'Speaker's Reference and Semantic Reference', *Midwest Studies in Philosophy* 2: 255–76. Page references to the reprint in P. French, T. Uehling, and H. Wettstein (eds.), *Contemporary Perspectives in the Philosophy of Language* (Minneapolis: University of Minnesota Press, 1979), 6–27.

Kvanvig, Jonathan L. 2007. 'Contextualism, Contrastivism, Relevant Alternatives, and Closure', *Philosophical Studies* 134: 131–40.

Kvart, Igal 1982. 'Quine and Modalities De Re: A Way Out?', *Journal of Philosophy* 79: 295–328.

Lewis, David 1973. *Counterfactuals*, Cambridge, Mass.: Harvard University Press.

——— 1976. 'Probabilities of Conditional and Conditional Probabilities', *Philosophical Review* 85: 297–315.

——— 1979. 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* 8: 339–59.

——— 1996. 'Elusive Knowledge', *Australasian Journal of Philosophy* 74: 549–67.

——— 2000. 'Causation as Influence', *Journal of Philosophy*, 182–97.

Ludlow, Peter 2005. 'Contextualism and the New Linguistic Turn in Epistemology', in Gerhard Preyer and Georg Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning and Truth*, Oxford: Oxford University Press, 11–50.

MacFarlane, John 2005. 'The Assessment Sensitivity of Knowledge Attributions', *Oxford Studies in Epistemology* 1: 197–233.

——— 2007. 'Relativism and Disagreement', *Philosophical Studies* 132: 17–31.

Malcolm, Norman 1952. 'Knowledge and Belief', *Mind* 51: 178–89.

——— 1976. 'Moore and Wittgenstein on the Sense of "I Know"', *Acta Philosophica Fennica* 28: 217–40. Reprinted (slightly revised) in N. Malcolm, *Thought and Knowledge*, Ithaca, NY: Cornell University Press, 1977.

Millikan, Ruth Garrett 1984. 'Naturalist Reflections on Knowledge', *Pacific Philosophical Quarterly* 65: 315–34.

Moore, G. E. 1959. *Philosophical Papers*, London: George Allen & Unwin Ltd.

———— 1962. *Commonplace Book: 1919–1953*, London: Allen & Unwin.

Nagel, Jennifer 2007. 'Epistemic Intuitions', *Philosophy Compass* 2/6: 792–819.

———— 2008. 'Knowledge Ascriptions and the Psychological Consequences of Changing Stakes', *Australasian Journal of Philosophy* 86: 279–94.

———— n.d. 'Knowledge Ascriptions and the Psychological Consequences of Thinking about Error', manuscript.

Plantinga, Alvin 1995. 'Précis of Warrant: The Current Debate and Warrant and Proper Function', *Philosophy and Phenomenological Research* 55: 393–6.

Quine, Willard Van Orman 1977. 'Intensions Revisited', *Midwest Studies in Philosophy* 2: 5–11. Page references to the reprint in P. French, T. Uehling, and H. Wettstein (eds.), *Contemporary Perspectives in the Philosophy of Language* (Minneapolis: University of Minnesota Press, 1979), 268–74.

Richard, Mark 2004. 'Contextualism and Relativism', *Philosophical Studies* 119: 215–42.

Rysiew, Patrick 2001. 'The Context-Sensitivity of Knowledge Attributions', *Noûs* 35: 477–514.

———— 2005. 'Contesting Contextualism', *Grazer Philosophische Studien* 69: 51–69.

Salmon, Nathan 1987. 'How to Measure the Standard Metre', *Proceedings of the Aristotelian Society* 88: 193–217.

Sartwell, Crispin 1991. 'Knowledge is Merely True Belief', *American Philosophical Quarterly* 28: 157–65.

———— 1992. 'Why Knowledge Is Merely True Belief', *Journal of Philosophy* 89: 167–80.

Schaffer, Jonathan 2004. 'From Contextualism to Contrastivism', *Philosophical Studies* 119: 73–103.

———— 2005. 'Contrastive Knowledge', *Oxford Studies in Epistemology* 1: 235–71.

Schiffer, Stephen 1979. 'Naming and Knowing', *Midwest Studies in Philosophy* 2: 28–41.

———— 1995. 'Descriptions, Indexicals, and Belief Reports: Some Dilemmas (But Not the Ones You Expect)', *Mind* 104: 107–31.

———— 1996. 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society* 96: 317–33.

Slote, Michael 1979. 'Assertion and Belief', in J. Dancy (ed.), *Papers on Language and Logic*, Keele: Keele University Library.

Slovic, Paul, Fischhoff, Baruch, and Lichtenstein, Sarah 1982. 'Fact Versus Fears: Understanding Perceived Risk', in D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment under Uncertainty: Heuristics and Biases*, Cambridge: Cambridge University Press.

Sosa, David forthcoming. 'Dubious Assertions', *Philosophical Studies*.

Stalnaker, Robert 1975. 'Indicative Conditionals', *Philosophia* 5: 269–86.

Stanley, Jason 2002. 'Modality and What is Said', *Philosophical Perspectives* 16: 321–44.

———— 2004. 'On the Linguistic Basis for Contextualism', *Philosophical Studies* 119: 119–46.

———— 2005. *Knowledge and Practical Interests*, Oxford: Oxford University Press.

Stine, Gail 1976. 'Skepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies* 29: 249–61.

Stone, Jim 2007. 'Contextualism and Warranted Assertion', *Pacific Philosophical Quarterly* 88: 92–113.

Stroud, Barry 1984. *The Significance of Philosophical Scepticism*, Oxford: Oxford University Press.

Unger, Peter 1975. *Ignorance: A Case for Scepticism*, Oxford: Oxford University Press.

—— 1984. *Philosophical Relativity*, Minneapolis: University of Minnesota Press.

—— 1986. 'The Cone Model of Knowledge', *Philosophical Topics* 14: 125–78.

—— 1995. 'Contextual Analysis in Ethics', *Philosophy and Phenomenological Research* 55: 1–26.

Vlastos, Gregory 1985. 'Socrates' Disavowal of Knowledge', *Philosophical Quarterly* 35: 1–31.

Weiner, Matthew 2005. 'Must We Know What We Say?', *Philosophical Review* 114: 227–51.

White, Alan 1975. *Modal Thinking*, Ithaca, NY: Cornell University Press.

Williamson, Timothy 1996. 'Knowing and Asserting', *Philosophical Review* 105: 489–523.

—— 2000. *Knowledge and its Limits*, Oxford: Oxford University Press.

—— 2005a. 'Knowledge and Scepticism', in F. Jackson and M. Smith (eds.), *The Oxford Handbook of Contemporary Philosophy*, Oxford: Oxford University Press.

—— 2005b. 'Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge', *Philosophical Quarterly* 55: 213–35.

Yourgrau, Palle 1983. 'Knowledge and Relevant Alternatives', *Synthese* 55: 175–90.

This page intentionally left blank

Index